Topos Cosmology Beyond Dark Matter: Gluing Obstructions, Contextual Gravity, and the End of Symmetry as Assumption

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

January 7, 2026

Modern cosmology is built from reconstructions, not direct views: we infer global structure from limited access along the past light cone, then enforce symmetry and consistency to make the inverse problem solvable. This strategy has produced extraordinary successes, yet it also concentrates unresolved tensions: a dark sector whose ontology remains unknown, a persistent mismatch between quantum theory and general relativity, and an uneasy dependence on principles that may reflect human preferences for compressibility rather than fundamental structure. This paper proposes an alternative stance: treat the Universe as topos-like, meaning that "reality" is not a single global model imposed from above, but a family of local rule-contexts with explicit transition maps and gluing conditions. In this framework, symmetry is not a metaphysical axiom but a testable regularizer; consistency is not presumed but earned via descent. Dark matter can be reframed as a measurable gluing obstruction: the failure of locally inferred gravitational potentials from distinct probes to assemble into a global section under allowed translations. We develop a compact conceptual architecture, define the minimal mathematics needed for a physics-grade proposal, and lay out falsifiable signatures that distinguish "obstruction cosmology" from particle dark matter and from ad hoc varying-G fixes.

Keywords: topos cosmology, dark matter, gluing obstruction, contextual gravity, sheaf theory, descent, cohomology, weak lensing, galaxy rotation curves, cluster mergers, modified gravity, quantum nonlocality, Kochen-Specker, Bell inequalities, relational physics, multi-scale functors, renormalization, cosmological principle, symmetry breaking, inverse problems, causal structure, past light cone, epistemic regularization, model selection, emergence.

A collaboration with GPT-5.2-Thinking. 60 pages. CC4.0.

Outline

1. The Problem Statement: When Successful Models Become Epistemically Brittle
1.1 Reconstruction from a measure-zero surface and the inverse-problem trap
1.2 Symmetry and consistency as regularizers: when enforcement masquerades as discovery
1.3 The accumulating fault lines: dark sector ontology, GR/QM mismatch, nonlocal correlations
1.4 What a "new model" must accomplish: predictive constraint without hidden priors

2. The Core Proposal: The Universe as a Topos-Like Family of Rule-Contexts
2.1 Contexts as the primitive: instruments, scales, causal access, coarse-grainings
2.2 Objects and morphisms: what counts as a state, an observable, and a translation
2.3 Internal logic and truth: why "what is real" becomes context-indexed
2.4 Global structure as earned gluing, not assumed geometry

3. Consistency Recast: Descent as a Physical Criterion
3.1 The sheaf condition as an empirical demand: agreement on overlaps
3.2 Global sections vs local sections: what "a single Universe model" really means
3.3 Failure modes: controlled inconsistency vs contradiction vs underdetermination
3.4 How descent replaces informal appeals to "self-consistency"

4. Symmetry Recast: Natural Transformations Instead of Metaphysical Invariances
4.1 Symmetry as invariance under context-change maps
4.2 Approximate symmetry, emergent symmetry, and scale-dependent symmetry
4.3 Parameterizing symmetry breaking: turning priors into estimable quantities
4.4 A discipline against overfitting: symmetry as compression with a price

5. Dark Matter as a Gluing Obstruction
5.1 The sheaf of gravitational inferences: dynamics, lensing, gas, time delays
5.2 Obstruction classes: when potentials cannot be glued into a global section
5.3 Interpreting "halos" as witnesses of failed descent
5.4 Why this avoids reifying bookkeeping terms into ontological substances

6. Contextual Gravity: Gravity as Curvature of Translation Rules
6.1 Connections on inference: transporting clocks, rods, and energies across contexts
6.2 Curvature as path dependence: non-commuting translations as "gravitational signal"
6.3 Recovering GR as a limit case: when translations approximate a smooth manifold
6.4 Where deviations should appear first: outskirts, mergers, voids, multi-probe seams

7. Quantum Nonlocality Reframed: No Global Section Instead of "Spooky Influence"
7.1 Measurement contexts as opens: outcomes as local sections
7.2 Bell and Kochen–Specker as gluing impossibilities

7.3 Preserving causality without classical realism: signaling vs correlation
7.4 A bridge concept: contextuality and dark-sector obstructions as the same kind of failure

8. Multi-Scale Topoi: Renormalization as Functorial Flow
8.1 A topos tower indexed by scale: coarse-graining as a functor
8.2 Universality as stable structure under functorial descent
8.3 Dark sector as an anomaly of commuting diagrams: the “counterterm” analogy
8.4 GR/QM mismatch as incompatible internal logics connected by scale functors

9. Falsifiability: What This Framework Must Predict (and Risk)
9.1 Cross-probe residual structure: mandatory relations between lensing and dynamics
9.2 Cluster collisions as separation tests: where obstructions must localize
9.3 Null predictions: when the obstruction must vanish (and why that matters)
9.4 Model selection penalties: degrees of freedom and the “no free function per object” rule
9.5 Distinguishing three programs: particle dark matter, modified gravity, obstruction cosmology

10. A Minimal Research Program and Prototype Implementation
10.1 Define the context category: a practical first choice for astronomy
10.2 Data objects and morphisms: what to store, what to map, what to compare
10.3 Computing obstructions from real datasets: a step-by-step pipeline
10.4 Benchmarks and shared tasks: rotation curves + weak lensing + cluster maps
10.5 Failure cases and refinement: when the framework collapses or becomes trivial

11. Implications and Limits: What Changes, What Stays
11.1 The cosmological principle as a convenience prior, not a foundation
11.2 Realism without omniscience: global claims as derived, not assumed
11.3 Why this does not “solve everything”: humility about ontology and ultimate causes
11.4 What success would look like: fewer knobs, more cross-context predictive power

12. Conclusion: From Enforced Elegance to Earned Structure
12.1 Restating the pivot: models as glued rule-contexts
12.2 Dark matter reinterpreted as obstruction, not entity
12.3 The new criterion: functorial predictability across contexts
12.4 Next steps: the smallest decisive tests

References

Art

1. The Problem Statement: When Successful Models Become Epistemically Brittle

Cosmology is unusually exposed to a recurring hazard in the philosophy of science: a model can be both empirically successful and epistemically fragile. The success comes from compressing vast, noisy, indirect observations into a small set of stable parameters and equations. The fragility comes from the fact that the compression depends on priors that are rarely optional in practice: assumptions about symmetry, consistency, and the legitimacy of extending local physics into a global narrative. When those priors are strong, the model can become "brittle" in a specific sense: it retains predictive power inside the domain where its assumptions are tacitly enforced, but it becomes difficult to diagnose whether the predictions are earned by nature or by the scaffolding. This brittleness is not a moral failure of cosmology; it is a structural consequence of attempting global inference from limited causal access. The proposal of this paper begins exactly here: if the inference problem is underdetermined, then the first responsibility is to expose which constraints are empirical and which are epistemic regularizers, and then to ask whether the regularizers are being mistaken for discoveries.

1.1 Reconstruction from a measure-zero surface and the inverse-problem trap

At any given moment, our astronomical observations arrive along the past light cone: a null hypersurface in spacetime. In a strict geometric sense, that surface is "thin" relative to the full 4D history it samples; it does not provide omniscient access to the interior. Even when we enrich the data by moving forward in time and observing repeatedly, our access remains causally structured: we do not observe "the Universe" in the way one observes a laboratory apparatus from multiple unconstrained viewpoints. We observe a stream of photons (and a smaller stream of neutrinos, cosmic rays, and gravitational waves) whose provenance is inferred, not directly tagged.

This causal limitation turns cosmology into an inverse problem. The forward problem is easy to state: given a global model (geometry, matter content, evolution laws), predict what a worldline-observer would see on their past light cone. The inverse problem is the hard one: given light-cone data, infer the global model. Inverse problems are generically underdetermined: distinct internal realities can project to similar observational data, especially when you allow flexible source modeling, selection effects, and astrophysical complexity. The trap appears when a reconstruction pipeline becomes so standardized that it feels like direct perception. Then "the inferred interior" is unconsciously treated as "the observed exterior," and the boundary between data and model disappears. A model can thereby gain rhetorical solidity without a commensurate increase in logical necessity.

The core implication is not skepticism for its own sake. It is that cosmology's basic epistemic posture should be: global statements are not first-order observations; they are second-order

reconstructions. That posture motivates the central move of this paper: represent the inferential structure explicitly, rather than hiding it behind a single global ansatz.

1.2 Symmetry and consistency as regularizers: when enforcement masquerades as discovery

Symmetry and consistency are the two most powerful tools for making an underdetermined inverse problem tractable. They are also the two easiest tools to confuse with "what the Universe is." In practice, cosmology often proceeds by selecting a symmetry class first, then estimating parameters inside that class. Homogeneity and isotropy are not merely descriptive; they drastically reduce degrees of freedom. Likewise, "consistency" assumptions—one cosmic time slicing, one global metric family, one matter model that can be pushed across scales—are not merely virtues; they are what make the reconstruction algorithm converge.

In statistical language, symmetry and consistency function as regularizers. They penalize complexity, suppress noise-fitting, and enforce reuse: the same rule must explain many observations at once. That is a legitimate scientific strategy. The hazard is the quiet inversion of logic: because the regularizers enable stable parameter inference, the resulting stability is then read as evidence that the regularizers were "discovered" rather than "imposed." A model's elegance becomes a psychological amplifier: the very property that makes the inference feasible becomes a reason to believe the model is fundamentally true.

This is not a call to abandon symmetry. It is a call to demote symmetry from axiom to hypothesis. A symmetry assumption should be treated like any other modeling choice: introduce explicit symmetry-breaking degrees of freedom, measure whether the data actually demand them, and quantify how much predictive power is lost or gained when the symmetry is softened. When symmetry is not parameterized, it cannot be tested; it is only enforced. Likewise, "consistency" should be formalized as a gluing criterion rather than asserted as a default. If local descriptions do not glue, the failure should not be forced into the model as "new stuff"; it should be diagnosed as a failure of translation rules or of the assumption that a single global section exists.

1.3 The accumulating fault lines: dark sector ontology, GR/QM mismatch, nonlocal correlations

The current scientific landscape contains a pattern that is difficult to ignore: several foundational issues are not independent; they cluster around the same epistemic seams.

First, there is the dark sector. The gravitational residuals that motivate dark matter and dark energy can be organized into highly predictive phenomenology. But the ontology remains unresolved: the "what is it?" question persists despite decades of effort. This is not a small gap. It is a sign that we may have a stable effective description without a settled underlying entity

description. In many sciences, that is a clue to re-check whether a bookkeeping construct has been prematurely reified.

Second, there is the mismatch between general relativity and quantum theory. Each is spectacularly successful within its domain; their cohabitation is conceptually uneasy. Cosmology, because it deals with the whole Universe and extreme regimes, is the arena where that mismatch becomes impossible to ignore. The absence of a clean unification is not merely technical; it suggests we may be using two different internal logics to describe nature, with no explicit account of how one logic becomes the other across contexts.

Third, there are the "counter-intuitive" features of quantum correlations. Whether one prefers to describe these as nonlocality, contextuality, or the impossibility of global classical assignments, the phenomenon has a common structure: local descriptions are consistent within their contexts, but there is no single global classical picture that glues them all without contradiction. That structural motif is precisely the motif that a topos-like approach is designed to formalize.

These fault lines are "adding up" in the sense that they all point to the same diagnosis: global, single-logic, single-section pictures may be over-asserted relative to the nature of our access and the structure of the phenomena. If the Universe behaves like a patchwork of contexts with constrained translations, then "unknown components" and "counterintuitive correlations" may be symptoms of insisting on a global section that does not exist in the naive category we are using.

1.4 What a "new model" must accomplish: predictive constraint without hidden priors

A new model worth taking seriously cannot be only a new vocabulary. It must do what successful models do: constrain prediction. But it must do so without simply smuggling in the old priors under new names. That sets a high bar, and it clarifies what "out-of-the-box reasoning" must mean here.

First, the model must make the inferential structure explicit. It must separate data objects (what is observed) from translation rules (how observations are related across contexts) and from global claims (what is inferred). If the model presumes that global claims exist, it must say exactly what gluing conditions guarantee that existence.

Second, the model must reduce degrees of freedom relative to ad hoc alternatives. In particular, it must forbid the temptation that plagues many "variable G" or "modified gravity" responses: a free function per galaxy. If the model introduces flexibility, that flexibility must be functorial: a single rule that applies across many contexts, not a per-object knob set.

Third, the model must produce cross-probe constraints. It is not enough to fit dynamics or lensing in isolation. A genuinely new framework should predict relationships between different

observational channels because those channels are local views of the same underlying rule-structure. If the model cannot force such relationships, it is not a physical theory; it is a re-description.

Fourth, the model must risk being wrong in a decisive way. It should have clear null predictions: circumstances under which the "obstruction" must vanish, or under which symmetry-breaking terms must be zero, or under which translation curvature must commute. Without such risks, the model becomes unfalsifiable philosophy.

Finally, the model must be expandable into a practical research program. It should tell us what to compute, from which datasets, with what pipeline, and what outcome would count as success or failure. The topos-like stance becomes physics only when it yields operational criteria for gluing, measurable obstruction classes, and predictive signatures that can be checked against data.

This section therefore sets the framing for the rest of the paper: the goal is not to discard successful cosmology, but to identify where its success rests on enforced regularities, then propose a replacement architecture in which regularities are tested as gluing outcomes rather than assumed as foundations.

2. The Core Proposal: The Universe as a Topos-Like Family of Rule-Contexts

The central move of this paper is to replace the default cosmological posture—one global spacetime model plus matter fields—with a different primitive: a structured family of contexts. Each context is a regime in which we can make and test statements, gather data, and apply a local rulebook. The Universe, on this view, is not first given as a global section; it is reconstructed (if possible) by gluing local descriptions along principled translations. "Topos-like" here signals a specific kind of architecture: reality is represented not by a single set of propositions evaluated in classical logic, but by a network of contexts with an internal logic of truth that may be weaker than global classical truth. The purpose is not to deny the possibility of a global model; it is to stop presuming it. A global model becomes a derived object—an earned compression—rather than an axiom.

This shift does two things at once. It makes explicit the inferential scaffolding that cosmology often hides, and it provides a formal place for persistent mismatches (dark-sector residuals, cross-probe tensions, quantum contextuality) to appear as structured obstructions rather than as metaphysical embarrassments. The remainder of this section defines the minimal vocabulary required to treat cosmology as a topos-like enterprise, while remaining close enough to standard practice that the proposal can be operationalized.

2.1 Contexts as the primitive: instruments, scales, causal access, coarse-grainings

A “context” is any situation in which the questions we are allowed to ask and the data we are allowed to collect are jointly constrained by physical access and by methodological choice. Cosmology is not one context; it is a patchwork of contexts glued by inference. A topos-like approach begins by admitting that patchwork explicitly.

Contexts can be indexed in several overlapping ways:

First, by instrument and measurement modality. Optical photometry, spectroscopy, radio interferometry, X-ray imaging, gravitational lensing reconstruction, gravitational-wave inference, and particle detections each define a context with different systematic errors, different forward models, and different “natural observables.” The “same object” viewed in two modalities is not trivially the same data; it is the output of different measurement functors applied to a shared but unobserved interior.

Second, by scale. The rules that describe laboratory quantum systems, solar system dynamics, galaxy rotation curves, cluster lensing, and horizon-scale structure are not automaticallys just the same rules with bigger numbers. They involve different approximations, different dominant degrees of freedom, and different coarse-grainings. Even if one believes there exists a single underlying law, our access to it is mediated by scale-dependent effective descriptions. A context therefore includes an explicit statement of the scale regime and the effective theory presumed valid there.

Third, by causal access. Observational cosmology is constrained by what can reach us: the past light cone, selection effects, and the finite age and opacity of the Universe. But even within that, causal access differs by redshift, by line-of-sight obstruction, by astrophysical foregrounds, and by horizon limits for certain messengers. A context includes which events are causally connectable to the observer and which are not, because that determines what “global” claims are even meaningful.

Fourth, by coarse-graining choice. Every inference pipeline in cosmology performs coarse-graining: it compresses data into summary statistics, templates, parametric fits, or latent-variable models. Choosing a power spectrum, a correlation function, a halo model, a profile family, or an equation-of-state parameterization is not an innocent technical choice; it is a choice of which information is preserved and which is discarded. In a topos-like stance, coarse-graining is part of the context, not a hidden step.

The key idea is that contexts are not subjective whims. They are structured by real physical constraints and by the practical necessity of doing science with finite data. But treating them as primitive prevents a common confusion: the confusion between a locally valid, instrument-limited, coarse-grained description and a globally valid ontology.

2.2 Objects and morphisms: what counts as a state, an observable, and a translation

A topos-like cosmology requires two explicit ingredients: the "things" that live in each context and the "maps" that relate contexts.

Within a context, the basic "objects" can be understood as structured collections of data and rules. A minimal object might include:

- A set of admissible system descriptions (candidate states) consistent with that context's measurement constraints.
- A set of observables, understood not as metaphysical properties but as the functions of the world that the context can operationally extract.
- A forward model that maps candidate states to predicted measurement outcomes within the context.

In ordinary physics language, we often speak as if there is one state of the system and one set of observables. In a context-first view, a "state" is always at least partially epistemic: it is the equivalence class of underlying possibilities that remain after we restrict to what the context can discriminate. That does not make it "mere belief"; it makes it a context-indexed representation of what is knowable and testable.

The second ingredient is morphisms: translation rules between contexts. A morphism is not merely a coordinate transform; it is a map that tells you how to move information, predictions, and uncertainties from one context to another. Examples include:

- Transformations between measurement modalities (e.g., from a mass distribution hypothesis to lensing shear predictions, or to dynamical velocity dispersion predictions).
- Coarse-graining maps between scales (e.g., from microphysical degrees of freedom to effective fluid parameters).
- Observer-change maps (e.g., relating inferences made by different observers or at different epochs).
- Model-to-model maps (e.g., approximations and limits, such as a relativistic model reducing to a Newtonian regime).

The point of morphisms is to formalize what cosmology already does informally. In standard practice, we routinely "translate" between probes: a mass model is fit to dynamics, then checked against lensing; a parameter inferred from one dataset is carried into another. These translations are where hidden priors often live. A topos-like framework forces them into the open: if the translation is not explicit, it is not allowed to count as evidence of global consistency.

This also provides a clean place for disciplined skepticism. If two contexts disagree, the disagreement is not immediately an existential crisis. It is a localized fact: the proposed morphisms do not commute, or the overlap is mis-modeled, or the object is ill-defined. That is already progress: the failure becomes an analyzable structural property rather than a foggy "tension."

2.3 Internal logic and truth: why "what is real" becomes context-indexed

Once contexts are primitive, classical global truth becomes optional. This is not a rejection of realism; it is a refusal to conflate realism with omniscience. In a topos-like setting, each context has an internal logic: a way to assign truth-values to propositions based on what can be verified or falsified within that context.

In ordinary classical physics, we often treat propositions as having determinate truth-values independent of measurement context. Quantum theory already strains that posture: measurement contexts do not always permit a single global assignment of properties. A topos-like approach generalizes this lesson beyond quantum foundations into cosmology: even classically, many cosmological propositions are not directly decidable from any single observational context. Their truth, to the extent it is meaningful, is mediated by the existence of reliable translations and gluing.

Thus "truth" becomes a statement about stable inference under allowed morphisms. A proposition counts as robust if it remains true across all contexts in which it can be transported and checked. This defines a notion of objectivity that does not require a God's-eye view. It also defines a principled humility: if a claim cannot be transported, it is not necessarily false, but it is not licensed as a global fact.

This has a direct payoff for the psychological critique. If symmetry and consistency are cognitive regularizers, then the question is: do they persist as stable truths under context-change, or do they only look true because we have restricted ourselves to contexts where they are hard-wired? A topos-like approach makes this question formal.

2.4 Global structure as earned gluing, not assumed geometry

The final move is the most important. Standard cosmology begins with a global geometric ansatz and then fits it. A topos-like cosmology begins with local rule-contexts and then asks whether they glue into a global structure.

In practical terms, "gluing" means the following: if two contexts overlap (they apply to the same physical systems in some regime), then the objects they assign and the predictions they make should agree after translation along the appropriate morphisms. If you have a cover of contexts that collectively span your domain, then consistent agreement on overlaps implies the existence of a global object (a global section) that restricts to each local description.

This shifts the meaning of "global geometry." Geometry becomes one possible global section—a particularly powerful compression—rather than a prior. If the gluing works, you may recover something that looks like a manifold with a metric and matter fields, and you may recover standard GR as an effective description of the commutativity of certain translations. But if the gluing fails in structured ways, the failure is not automatically patched by inventing new substances. The failure is diagnosed as a gluing obstruction: the contexts cannot be jointly satisfied under the proposed translation rules.

This is the conceptual place where dark matter can be reframed. If dynamics and lensing can be made locally consistent in their own contexts, but cannot be made jointly consistent as restrictions of a single global gravitational potential under the allowed morphisms, then the "extra mass" inference is a symptom of failed descent. The "halo" is not an entity one must believe in; it is a witness that gluing failed under the assumed translation rules. The correct scientific next step is then to ask which translation rule is wrong, which context is mis-modeled, or whether the global-section assumption itself should be replaced by a weaker global object (a stack-like structure, a torsor-like twist, an obstruction class).

The price of this move is that it must eventually touch data. A topos-like cosmology is not permitted to remain at the level of metaphors. Its advantage is that it naturally supplies what cosmology most needs: a formal language for tracking what is enforced, what is transported, what is glued, and where and how the gluing fails. In the sections that follow, we will treat descent and symmetry not as virtues but as measurable outcomes, and we will define what it would mean for "obstruction cosmology" to be predictive rather than merely evocative.

3. Consistency Recast: Descent as a Physical Criterion

"Consistency" is often invoked in cosmology as if it were a moral virtue of good theories or a metaphysical requirement of the Universe. In practice, it functions as an inference rule: we expect that independently constructed descriptions should agree where they overlap, and we use that expectation to eliminate degrees of freedom. The problem is not that this expectation is wrong; it is that it is rarely made explicit. When "consistency" is implicit, it cannot be tested, relaxed, or localized. A topos-like approach turns consistency into a concrete physical criterion: descent. Descent is the statement that local descriptions can be glued into a coherent global description by explicit transition maps, and that the glued object reproduces the locals when restricted back down. In this section, "descent" is not treated as abstract category theory jargon, but as the exact thing cosmology claims whenever it asserts that one model explains multiple datasets simultaneously. The advantage of the descent framing is that it distinguishes three very different situations that are often conflated: genuine inconsistency (a contradiction), controlled inconsistency (a principled mismatch that encodes missing structure), and

underdetermination (multiple global candidates fit the same local data). These are epistemically distinct, and physics should respond differently to each.

3.1 The sheaf condition as an empirical demand: agreement on overlaps

A sheaf condition is, in plain terms, an agreement rule. Suppose you have a family of contexts that collectively cover a domain of inquiry. Each context yields a local description: a rulebook, a set of inferred quantities, or a candidate model class. The overlap between two contexts is the regime where both apply to the same systems, or where both can be compared through a translation pipeline. The sheaf condition says: if two local descriptions agree on overlaps (after applying the appropriate translations), then there exists a larger description that unifies them; and conversely, any unified description must restrict to the locals.

In cosmology, overlaps are everywhere, but they are rarely treated as first-class objects. "Overlap" can mean: two probes of the same mass distribution, two datasets that should share parameters, two regimes that should share a limiting behavior, or two descriptions connected by an effective-theory approximation. The descent stance forces an explicit question: what exactly is being compared, and what exact map makes them comparable?

This is not merely a philosophical demand. It is the empirical heart of multi-probe cosmology. When a model claims to explain both galaxy dynamics and gravitational lensing, it is asserting the existence of a single underlying structure whose projections into those probes are compatible. That is a sheaf-like claim. In ordinary practice, one often discovers that "agreement" depends on extra nuisance parameters, calibration assumptions, selection effects, or modeling choices about baryonic physics. A descent-first approach does not treat those as annoyances; it treats them as part of the overlap structure. If agreement only occurs after adding flexible "bridges," then the model has not earned unity; it has purchased it.

Thus the sheaf condition becomes an empirical demand with three components:

- Define the local objects (what each context outputs).
- Define the restriction and comparison maps (how locals are transported).
- Test whether the transported locals agree on overlaps within declared uncertainties.

Only then may one claim that a unified object exists. "Consistency" is no longer an adjective; it is a measurable predicate of a proposed gluing structure.

3.2 Global sections vs local sections: what "a single Universe model" really means

The phrase "a single Universe model" hides a technical assumption: that there exists a global section. A global section is a single, coherent assignment of the relevant structure across all contexts such that each context's local description is merely a restriction of the global one.

In everyday terms: a global section is what you are asserting when you say “there is one metric and one matter content and one set of parameters that jointly explain all these observations.” It is not merely the claim that you can fit each dataset; it is the stronger claim that the fits are compatible because they arise from one underlying object.

Local sections are weaker: they are context-valid descriptions that may work well in their own domains but do not automatically glue. A local section is what you get when you say “in this probe and this regime, the inference is stable,” without committing to global unification.

Cosmology often proceeds by acting as if the existence of a global section is the default. That default has tremendous advantages: it permits parameter estimation, model comparison, and tight forecasts. But it also has a subtle cost: when gluing fails, the failure is not interpreted as “perhaps there is no global section in this category,” but as “the global section exists, so the mismatch must be fixed by adding new components.” In other words, dark sectors can become the price of insisting on global sectionhood.

A descent-centered approach makes the global section a hypothesis, not a background assumption. It asks: do the local sections we actually infer from different contexts glue into a single object under declared translations? If they do, the global model is earned and strengthened. If they do not, we can respond in more than one way: revise the translations, refine the local objects, enlarge the category of global objects (e.g., allow twisted sections), or accept that the best description is inherently local-with-structure rather than global-and-smooth.

This is the technical meaning behind the paper’s claim that global structure should be “earned gluing, not assumed geometry.” A global cosmology is not “what cosmology must be”; it is “what descent certifies.”

3.3 Failure modes: controlled inconsistency vs contradiction vs underdetermination

Once consistency is formalized as descent, “failure” stops being a single undifferentiated alarm bell. There are at least three distinct failure modes, and confusing them creates much of the conceptual heat in foundational debates.

First is contradiction. This is the hard failure: no matter how you transport and compare, the local descriptions cannot agree, and the disagreement cannot be attributed to declared uncertainties. Contradiction demands either that at least one local inference is wrong (systematics, mis-modeled forward physics) or that the translation maps are wrong (the rule linking contexts is incorrect). Contradiction is an empirical event, not a philosophical discomfort.

Second is underdetermination. Here, the locals do not uniquely determine a global object. Multiple globals restrict to the same locals (within uncertainty). This is extremely common in inverse problems: different internal structures can project to similar observational data.

Underdetermination is not inconsistency. It does not imply the model is false; it implies the evidence does not force uniqueness. The correct response is to seek new overlaps, new probes, or new invariants—not to declare metaphysical crisis. A descent framework makes underdetermination visible by exhibiting multiple candidate gluings.

Third is controlled inconsistency (or structured non-gluability). This is the most interesting case for the present proposal. Here, local descriptions are internally coherent within their contexts, but attempts to glue them fail in a way that is repeatable, localized, and patternful. The failure is not random noise; it has structure. In mathematics, such failures often correspond to obstructions: a nontrivial cohomology class, a torsor without a section, a twist that prevents global trivialization. In physics language, controlled inconsistency suggests missing structure in the translation rules or missing degrees of freedom that only appear in overlaps.

This is precisely the kind of failure that can be mistaken for "evidence of new substance." Dark matter, in the present framing, is a candidate instance: a structured mismatch between locally inferred gravitational content across probes and scales. The descent view does not pre-judge whether the correct completion is a particle, a field, a modified translation rule, or a change in what counts as a global object. It insists only that the mismatch be treated as a gluing diagnostic rather than automatically as ontology.

These three failure modes are not rhetorical niceties. They are different scientific situations with different prescriptions:

- Contradiction: locate and correct error or revise rule.
- Underdetermination: add data or constrain model class.
- Controlled inconsistency: identify the obstruction and classify completions.

3.4 How descent replaces informal appeals to "self-consistency"

"Self-consistency" is often used as a shield in theoretical physics: a model is said to be good because it is consistent with itself and with known constraints. But in cosmology, many "consistency" claims are not internal logical checks; they are claims of cross-context agreement. Without descent, those claims remain informal and therefore psychologically vulnerable. One can enforce them without noticing one is enforcing them.

Descent replaces informal self-consistency in four ways.

First, it forces explicit overlaps. Instead of saying "the model matches everything," one must specify which contexts overlap, what the overlap means operationally, and what quantity is compared.

Second, it forces explicit translations. Instead of importing a parameter value from one domain into another as if it were a direct observable, one must define the morphism that carries it and tracks how uncertainty and modeling assumptions transform.

Third, it produces explicit certificates and explicit obstructions. If descent holds, that fact can be stated as the existence of a glued object that restricts correctly everywhere tested. If descent fails, the failure can be localized: which overlaps fail, by how much, and with what pattern. This turns "tension" into structure.

Fourth, it changes the ontology reflex. Instead of responding to every mismatch by adding an entity, descent asks first whether the mismatch is telling you that you attempted to glue in the wrong category. Sometimes the right response is indeed a new physical component. But sometimes the right response is a new translation rule, a new effective description, or a recognition that global sectionhood was an unjustified demand.

In short, descent turns "consistency" from a vibe into a criterion. It converts philosophical discomfort into an operational pipeline: define contexts, define overlaps, define morphisms, test commutativity, and classify failures. That pipeline is what makes a topos-like cosmology more than a metaphor. It provides the scaffolding for the paper's next claims: symmetry is a testable compression rule, dark matter can be an obstruction class, and "global cosmology" is a derived object whose existence must be demonstrated rather than assumed.

4. Symmetry Recast: Natural Transformations Instead of Metaphysical Invariances

Symmetry occupies a privileged place in physics because it is an engine of compression. It allows one rule to stand in for many cases: rotate the apparatus, translate the origin, change coordinates, or relabel equivalent degrees of freedom, and the predictions remain invariant. That compression is not merely aesthetic; it is what makes generalization possible when data are sparse relative to the space of possible worlds. The trouble arises when symmetry is treated as metaphysics rather than method. In cosmology, symmetry is often elevated from a modeling choice into a foundational claim about the Universe "in itself." This elevation is psychologically appealing because it stabilizes inference: if you assume homogeneity and isotropy, the inverse problem becomes tractable. But tractability is not the same as truth.

A topos-like posture proposes a different definition of symmetry. Symmetry is not a proclamation about reality "out there." Symmetry is a statement about the stability of predictions under allowable changes of context. It is a property of the translation structure: if you move a description from one context to another by a legitimate map, the description's content is preserved in a precise way. This is exactly the idea of a natural transformation: the claim that a diagram commutes, that transporting then measuring is equivalent to measuring

then transporting. Symmetry becomes something you can test and localize. It is no longer a metaphysical guarantee; it is a measured regularity of how contexts relate.

4.1 Symmetry as invariance under context-change maps

In a context-first framework, the basic question is not "what symmetries does the Universe possess?" but "which context changes leave our predictions invariant?" A context change might be as simple as a coordinate transformation, or as complicated as shifting instrument modality, redshift regime, or coarse-graining procedure. If we can define a map that legitimately carries information from one context to another, then symmetry is the statement that the transported description agrees with the native description in the target context.

This reframing has two consequences.

First, it makes symmetry operational. Instead of a slogan, symmetry becomes a commutativity condition: take a local object in one context, transport it through the morphism, and compare. If the outputs match across all relevant observables within uncertainty, the symmetry holds for that pair of contexts. If not, the symmetry fails, and the failure is localized.

Second, it expands what counts as symmetry. Traditional symmetries (rotations, translations, Lorentz invariance) are special cases of a more general idea: invariance under changes of description that should not matter. In cosmology, many such changes are not geometric but inferential: choice of tracer population, choice of bias model, choice of lensing reconstruction method, choice of smoothing scale, and choice of parametric family. When these choices are treated as "details," symmetry can be smuggled in by assuming that different choices yield equivalent inferences. When contexts are explicit, those equivalences become symmetry claims, and they can be tested rather than presumed.

This is the point at which "symmetry as a projection of the human mind" becomes a scientific program. We admit that symmetry is a compression preference, then we ask which compressions are stable under context change and which are not. The Universe gets to answer.

4.2 Approximate symmetry, emergent symmetry, and scale-dependent symmetry

Once symmetry is defined as a stability property under context change, it becomes natural to treat symmetry as approximate, emergent, and scale-dependent without conceptual distress.

Approximate symmetry is the rule rather than the exception. In real data, invariances are rarely exact. The meaningful question is: how strong is the invariance, across which contexts, and with what residual structure? An approximate symmetry is an empirical regularity characterized by a small symmetry-breaking term whose magnitude and pattern can be measured.

Emergent symmetry is a statement about coarse-graining. Many systems display symmetries at large scales that are not present microscopically. In cosmology, this is particularly relevant

because we are always observing effective descriptions. Homogeneity and isotropy, if they hold, plausibly hold as emergent symmetries of large-scale statistics rather than as microscopic properties. The topos-like framing makes this explicit: a symmetry can hold in the coarse-grained contexts even if it does not hold in finer contexts, and this is not a contradiction; it is a statement about which morphisms erase which irregularities.

Scale-dependent symmetry follows immediately. The symmetry class appropriate to a homogeneous expanding background does not describe the lumpy Universe at galaxy scales, and even within galaxy scales the "symmetry" of a halo is at best approximate and often environment-dependent. Treating symmetry as global metaphysics forces one into awkward language: "the Universe is homogeneous" but "locally it isn't." Treating symmetry as context-indexed stability removes the awkwardness. The Universe is symmetric in the contexts where transport commutes with prediction to a chosen accuracy. Elsewhere, it is not, and that lack is informative.

This is a critical conceptual upgrade. It allows one to avoid the binary trap in which "either the cosmological principle is true" or "it is false." In a descent-and-context stance, the meaningful question is: in which contexts is it a good regularizer, and what are the measured symmetry-breaking residuals where it fails?

4.3 Parameterizing symmetry breaking: turning priors into estimable quantities

The most important methodological step in demoting symmetry from metaphysics to physics is to parameterize its breaking. If a symmetry is hard-coded, it cannot be tested. If symmetry breaking is allowed and penalized, then the data can decide how much symmetry is warranted.

Parameterizing symmetry breaking has several layers, from minimal to expansive.

At the minimal level, one introduces a small set of deformation parameters that quantify departures from the symmetric ansatz. For cosmological isotropy, this might mean allowing a preferred-direction component in certain summary statistics. For homogeneity, it might mean allowing scale-dependent inhomogeneous terms or local void/bulk-flow distortions within the mapping from redshift space to distance. For halo modeling, it might mean allowing triaxiality, substructure, and environment-dependent profiles rather than enforcing spherical universality.

At the next level, one allows symmetry breaking as a controlled function class rather than a single parameter, but with complexity penalties that prevent overfitting. The point is not to introduce an arbitrary function; it is to allow flexibility while keeping it costly, so that symmetry is retained unless the data truly demand the additional freedom.

At the most principled level, symmetry breaking is treated as a property of the morphisms between contexts. Instead of adding "an anisotropy term" to the global model, one asks whether the translation rule between two contexts is itself non-commuting in a structured way.

In that case, symmetry breaking is not a property of "the Universe" in isolation, but a property of how descriptions relate across contexts. That is often closer to what is actually happening in practice: disagreements between probes frequently trace back to how a latent quantity is translated from one observational pipeline to another.

This parameterization move is the heart of your psychological critique. Symmetry is a prior. Priors are unavoidable. The scientific question is whether the prior is quietly deciding the outcome. By making symmetry-breaking estimable, you measure how much the prior is doing. A symmetry that survives this procedure is not merely enforced; it is earned.

4.4 A discipline against overfitting: symmetry as compression with a price

One reason symmetry became foundational in physics is that it protects against the universal failure mode of inverse problems: the temptation to fit noise. A new framework cannot simply "relax symmetries" and declare victory. Relaxation without discipline is a fast path to non-falsifiable flexibility.

The descent-and-context approach therefore insists on a specific discipline: symmetry is compression with a price. The "price" can be implemented as explicit penalties on symmetry-breaking degrees of freedom, as model selection criteria, or as the requirement that any added flexibility must improve out-of-sample prediction across contexts rather than merely improving in-sample fit within one context.

This gives a clean test for whether a symmetry is merely psychological or physically grounded: if symmetry-breaking complexity improves one dataset's fit but degrades cross-context predictive agreement, then the "improvement" was an overfit. Conversely, if allowing controlled symmetry breaking reduces cross-probe tensions and yields better joint predictions with a modest increase in complexity, then the symmetry was too rigid. Either outcome is informative, and both are testable.

In the specific aims of this paper, this discipline also blocks the typical escape route of ad hoc modified gravity models: allowing an effectively arbitrary modification per system. If one can tune a function per galaxy, one has not explained anything; one has merely re-labeled a discrepancy. A topos-like model must therefore enforce functorial symmetry-breaking: one rule that predicts where and how symmetry breaks across contexts. The rule is judged by whether it reduces the total description length of the world while increasing cross-context predictive power.

This section therefore completes the conceptual pivot: symmetry is not a metaphysical invariance imposed to make the math possible. It is a testable commutativity property of context-change maps, expected to hold approximately and sometimes emergently, and made scientific by explicit symmetry-breaking parameterization under complexity penalties. In the

next section, this machinery will be applied to the most charged case: the dark matter problem, reframed as a gluing obstruction rather than as a substance postulate.

5. Dark Matter as a Gluing Obstruction

The phrase “dark matter” functions today as a remarkably efficient compression: it names a persistent gravitational residual that appears across many systems and many probes. The compression is useful, but it carries an ontological temptation. A residual can be treated as an entity, and once treated as an entity it attracts a default narrative: “there is extra stuff.” The purpose of this section is to replace that reflex with a more disciplined alternative: interpret the residual first as a failure of gluing. In a context-first, descent-driven cosmology, “dark matter” is not a substance claim. It is a statement about the inability to glue locally inferred gravitational structures into a single global section under the currently assumed translation rules. The residual is then not “mysterious mass” by default; it is an obstruction that points to a specific kind of missing structure: either the translation maps are wrong, the local objects are misidentified, or the category of global objects is too restrictive.

This reframing does not deny the possibility that the obstruction’s best completion is a new particle or field. It simply refuses to grant that completion automatically. It insists that the first-class empirical fact is the structured mismatch across contexts, and that the appropriate mathematical object to represent such mismatch is an obstruction class. This section shows how to build that object from the probes cosmology already uses.

5.1 The sheaf of gravitational inferences: dynamics, lensing, gas, time delays

Begin with the practical reality: “gravity” is not measured directly. What is measured are effects: motions, distortions, temperatures, pressures, and time delays. Each probe yields an inference about the underlying gravitational influence, but each inference is filtered through its own forward model and systematics. In a topos-like formulation, each probe defines a context, and each context yields a local section of a structured sheaf of gravitational inferences.

A minimal “sheaf of gravitational inferences” can be described as follows. For each context c, we assign an object S(c) that represents the set (or distribution) of gravitational structures compatible with the data in that context. The word “structure” is chosen deliberately. It could be a potential, a metric perturbation, a mass density proxy, or a pair of potentials, depending on the regime and formalism. The key is that it is whatever quantity the context can infer robustly.

The main contexts typically include:

First, dynamics. From rotation curves, velocity dispersions, and equilibrium modeling, one infers an effective gravitational acceleration field or potential gradient. This is a local inference tied to assumptions about equilibrium, tracer populations, anisotropies, and baryonic mass-to-light

conversion. In the sheaf language, S_dyn(c) is the family of potentials compatible with the kinematic data under those assumptions.

Second, gravitational lensing. From shear maps and strong-lensing arcs, one infers a projected mass distribution or, more fundamentally, a lensing potential integrated along the line of sight. Lensing is sensitive to how gravity deflects null geodesics; its inference pipeline has its own degeneracies and calibrations. In the sheaf language, S_lens(c) is the family of lensing potentials consistent with the distortions observed.

Third, gas physics in galaxies and clusters. Hot gas distributions inferred from X-ray or SZ observations provide pressure and temperature profiles. Under an assumption like hydrostatic equilibrium (or under more general dynamical modeling), one infers the gravitational potential required to confine the gas. This yields another local gravitational inference S_gas(c), often with large and structured systematics because non-thermal pressure support and disequilibrium can be significant.

Fourth, time-delay and timing observables. In certain systems, one can infer gravitational potentials from time delay between multiple lensed images of a variable source, or from propagation effects in pulsar timing or other signals. These are not always available, but when they are, they add an independent channel whose forward model is different enough to be valuable. Call this S_time(c).

All of these inferences target “the same underlying gravity” in the naive global-section picture. But in reality, they are context-indexed. Each gives a partial view of gravitational influence, not an unfiltered truth. The sheaf approach formalizes this situation rather than smoothing it away. A key part of the formalization is the existence of restriction maps: if a context c’ is a refinement of a context c (higher resolution, tighter selection, narrower regime), then a gravitational inference in c should restrict to one in c’. Conversely, coarse-graining maps transport inferences from fine to coarse contexts. The sheaf is the entire network of such local inferences and the maps between them.

5.2 Obstruction classes: when potentials cannot be glued into a global section

With the sheaf of local inferences in place, “dark matter” can be reframed as a statement about non-gluability. If the local gravitational inferences from different probes and scales were restrictions of one global gravitational structure, then there should exist a global section G such that, for each context c, the restriction of G to c lies in S(c), and on overlaps the restrictions agree after applying the appropriate translation maps.

When this fails in a structured way, we say there is an obstruction to gluing. In mathematical language, the obstruction may live in a cohomology group associated with the cover of

contexts; in physical language, it is the residual mismatch that cannot be eliminated by any global gravitational structure within the assumed category and translation rules.

This framing matters because it distinguishes two different phenomena that are easily conflated:

- Local success: each context has a good fit within its own assumptions.
- Global failure: there is no single gravitational structure that simultaneously accounts for all contexts when compared on overlaps.

Dark matter, in this proposal, names the second phenomenon. It is not “something extra”; it is “the failure of a would-be global gravitational section.”

The obstruction class can be conceptualized in several equivalent ways.

One way is diagrammatic. Consider a cycle of translations: from dynamics to lensing to gas to dynamics again, all within an overlap regime (say, a cluster). If the translations commute and are faithful, you should return to where you started. If you do not, the cycle has holonomy: you accumulate a mismatch. That mismatch is the obstruction. It is not random; it has a direction and magnitude in the space of inferred gravitational structures.

Another way is comparative. On overlaps, you can compute a difference between two locally inferred potentials after transporting them to a common representation. If these differences cannot be driven to zero for any choice of global G within the allowed model class, but instead stabilize as a repeatable residual pattern, that pattern is the obstruction. In the simplest case, it appears as a systematic “missing mass” profile. In more refined cases, it appears as probe-dependent and environment-dependent structure.

A third way is categorical. Sometimes the obstruction is not that the Universe lacks a global section, but that we have asked for the wrong kind of global object. The correct global object might be a twisted section, a torsor, or a stack-like object that has local trivializations but no global trivialization. In that case, “dark matter” is the signal that the global object is nontrivial. The obstruction class then classifies what kind of twist is present.

All three interpretations share the same discipline: the residual is treated as a structural feature of gluing, not as immediate evidence for a new substance.

5.3 Interpreting “halos” as witnesses of failed descent

In standard cosmology, a “halo” is a distribution of dark matter around galaxies and clusters that explains rotation curves, velocity dispersions, and lensing. In the present framework, a “halo” is reinterpreted as a witness: it is a local object introduced to force gluing to work under the global-section assumption.

This is a subtle but crucial difference. A witness is an auxiliary object that certifies the failure of a naive gluing attempt. It is what you add when you insist on a single global potential but your local inferences disagree. The halo becomes the compensator that makes the overlap conditions pass.

This witness interpretation explains several features of the dark matter story without settling ontology:

- Why halo profiles seem to have flexible families. A witness is shaped to satisfy overlap constraints under limited data; its flexibility reflects the underdetermination of the inverse problem.
- Why baryonic modeling matters so much. If the local inferences depend on baryonic assumptions, then the obstruction depends on them; a witness halo absorbs the residual left over after baryonic choices.
- Why certain systems become "crucial tests." Systems like cluster mergers act like severe overlap constraints: they bring multiple contexts into direct comparison under extreme conditions. In witness language, they are contexts where forcing gluing requires particularly structured compensators.

None of this denies the possibility that the witness corresponds to something real. A witness can be an entity. But the witness interpretation changes the logical order. First we observe that gluing fails. Then we introduce a witness object that restores gluing. Only then do we ask whether the witness is best understood as a physical substance, a modification of translation rules, or a sign that our category of global objects should be enlarged.

That shift is the antidote to premature reification. It turns "halo" from a metaphysical commitment into a provisional certificate of non-gluability.

5.4 Why this avoids reifying bookkeeping terms into ontological substances

The primary virtue of the obstruction framing is epistemic hygiene. It keeps a sharp boundary between phenomenology and ontology.

Phenomenology: there exist structured residuals when attempting to unify gravitational inferences across probes and scales. This is an empirical claim and can be quantified.
Ontology: the residual is caused by a new particle, a new field, a modified gravity law, a breakdown of assumed translation rules, or a deeper contextual structure. This is a theoretical completion problem.

Standard discourse often collapses phenomenology into ontology because the word "dark matter" names both. The topos-like approach separates them. "Dark matter" becomes

shorthand for an obstruction class in the sheaf of gravitational inferences. The obstruction is what is measured. The completion is what is hypothesized.

This separation also clarifies what counts as progress. Finding tighter particle constraints is progress, but so is finding that the obstruction has a probe-dependent structure that cannot be captured by any particle-like completion without additional assumptions. Demonstrating that the obstruction localizes in certain overlap regimes but vanishes in others is progress, because obstructions that vanish under specific conditions imply a classification and therefore a pathway to explanation. Showing that a proposed completion makes the translation diagrams commute with fewer degrees of freedom is progress, because it replaces a witness with a law.

Finally, this framing reduces the psychological load. If one feels that symmetry and consistency are partly cognitive projections, then insisting that the Universe must provide a globally consistent, symmetric section is itself a projection risk. The obstruction view allows one to say: perhaps our demand for a global section is too strong; perhaps the Universe is locally coherent and globally twisted; perhaps "extra components" are the shadow of that twist. The point is not to declare victory for any alternative. The point is to create a formal language in which failure is informative rather than humiliating.

In the next sections, this idea will be pushed further: gravity itself can be treated as curvature of translation rules between contexts, and quantum contextuality can be understood as a related no-global-section phenomenon. In that combined picture, the dark sector becomes not an isolated anomaly but a repeated structural motif: the Universe may be comprehensible not as a single global story, but as a network of locally valid rule-contexts whose gluing is nontrivial and therefore physically significant.

6. Contextual Gravity: Gravity as Curvature of Translation Rules

If dark matter is reframed as a gluing obstruction, the natural next question is: what is "gravity" in the architecture that produces that obstruction? In standard physics, gravity is represented as geometry (GR) or as a force (Newtonian limit). In a context-first, topos-like framework, gravity can be reinterpreted as something slightly different but more primitive: the structured way in which inferences transport across contexts. On this view, "gravity" is not first a field on a pre-assumed manifold; it is the rule that tells you how clocks, rods, energies, and trajectories inferred in one context are transported into another context, and what invariants survive those transports.

This reframing preserves what is best about GR (gravity as a universal organizer of relational measurements) while changing where universality is encoded. Universality becomes a constraint on translation maps: in the regimes where GR holds, translations commute in a way

that permits a smooth manifold description. In regimes where “dark” residuals appear, translations may fail to commute; the resulting path dependence is not just mathematical inconvenience, but the observable signature of curvature in the inference-connection. In this section, we develop that idea in a physics-grade way: gravity as connection, gravitational field strength as holonomy, GR as the near-flat or integrable limit, and deviations as concentrated at the seams where contexts overlap but do not perfectly glue.

6.1 Connections on inference: transporting clocks, rods, and energies across contexts

A connection is a rule for transport. In ordinary differential geometry, a connection tells you how to transport vectors along curves so you can compare them at different points. In contextual gravity, the analogous task is: how do you transport inferential objects between contexts so that comparisons are meaningful?

The key inferential objects that must be transported are the operational primitives of physics: clocks (time standards), rods (distance standards), and energies (calibrations that link motion to matter). In cosmology, these primitives are never “given”; they are reconstructed from layered measurement pipelines. Transport is therefore the silent engine of every cosmological claim.

To make this explicit, consider what must be transported when you say “this galaxy has mass M” or “this cluster potential is Phi(r).” You are transporting at least four kinds of structure:

First, calibration structure. A clock is not “time”; it is a physical process whose ticks are counted and compared. A rod is not “distance”; it is a standard derived from light travel, spectral lines, or geometric baselines. These standards must be transported across instruments, epochs, and environments. A connection on inference must specify how calibration moves across contexts without accumulating untracked distortions.

Second, model structure. Inference is not measurement. Each probe uses a forward model: velocities imply potential gradients under dynamical assumptions; lensing shear implies projected potentials under reconstruction assumptions; gas profiles imply confining potentials under thermodynamic assumptions. Transport must specify how an inferred object in one model language becomes an inferred object in another model language. This is not optional; it is the definition of multi-probe science.

Third, uncertainty structure. A transported estimate without transported uncertainty is not transport; it is propaganda. The connection must specify how uncertainty, bias, and degeneracy transform under context-change. Otherwise, apparent “consistency” can be purchased by silently widening error bars in one domain or narrowing them in another.

Fourth, equivalence structure. Contextual objects are defined up to equivalence relations: what counts as the “same” potential may differ by gauge freedoms, additive constants, or

unobservable components in a given probe. Transport must respect those equivalences, or else it will generate spurious contradictions.

In this picture, a “connection on inference” is not a single formula; it is a disciplined rule-set that tells you how to carry operational primitives and inferred gravitational structures from one context to the next. When such a connection exists and is well-behaved, it allows a global story: clocks, rods, energies, and potentials line up across contexts, and the Universe looks like a single coherent geometry. When it fails, the failure is not merely error; it is data about the limits of integrability of our translations.

6.2 Curvature as path dependence: non-commuting translations as “gravitational signal”
Curvature is what you detect when transport depends on the path you took. In geometry, you parallel transport around a loop and see whether you return to the same vector. In contextual gravity, you transport inferences around a loop of contexts and see whether you return to the same inferred structure.

The basic diagnostic is simple in concept. Choose a set of contexts that overlap on a real system (for example: dynamics, lensing, gas, time delays). Each context yields a local inference about “gravity.” Now choose translation maps that take one inference into the representational language of another. If the translations are consistent, then any two routes that start and end in the same context should agree. If they do not, the mismatch is holonomy, and holonomy is curvature.

This makes “gravity as curvature of translation rules” operational:

- A translation loop is a sequence of morphisms that returns you to the starting context.
- The loop produces a residual: what changed after you went around once.
- That residual is the contextual curvature signal.

Two features of this are important.

First, this curvature is not merely “mathematical.” It is measurable because it manifests as stable, probe-dependent residuals that do not vanish under improved statistics if the translations themselves are missing structure. In the standard narrative, such residuals are often treated as evidence for extra mass. In the contextual narrative, they are evidence that the translation rules are non-integrable in the chosen category of global objects.

Second, this curvature is naturally localized. It can be large at certain seams (for example, where equilibrium assumptions break, where projection effects dominate, where baryonic feedback reshapes distributions, where collisionless and collisional components separate) and small elsewhere. That localization is exactly what you would expect if “dark matter” is an obstruction

class: the obstruction is not necessarily uniformly spread; it is concentrated where gluing is hardest.

This provides a new way to interpret a familiar fact: different probes "weigh" systems differently. In the usual view, that is either systematic error or evidence for unseen mass. In contextual gravity, it is the direct signal of non-commuting translations. The question becomes: does there exist a single completion that makes the diagrams commute with minimal added structure? If yes, the curvature is explained. If no, the curvature is a fundamental feature of how contexts relate.

6.3 Recovering GR as a limit case: when translations approximate a smooth manifold

A topos-like approach is not a rejection of GR. It is a proposal about when and why GR is the right compression.

GR can be recovered as a limit case when the translation structure becomes effectively integrable. Integrability here means: there exists a representation in which the transport of clocks, rods, and energies is consistent across contexts, and path dependence becomes negligible (or becomes exactly the kind of path dependence encoded by a smooth connection on a manifold).

In practical terms, GR emerges when three conditions hold.

First, a stable notion of local inertial frames exists across contexts. This is equivalent to saying that local calibrations can be transported without ambiguity: clocks and rods defined in one local regime can be compared to those in another via translation maps that do not introduce context-specific distortions beyond what GR already accounts for.

Second, the translation maps commute sufficiently well that they can be summarized by a smooth geometry. When the loop residuals (holonomies) behave like the holonomies of a smooth connection, you can package them into curvature tensors and speak in the familiar language of spacetime curvature. The key point is that GR's geometry becomes the unique efficient language for a translation structure that is consistent enough.

Third, the choice of context becomes "thin." In many regimes of physics, the differences between probes can be treated as different windows onto the same underlying structure, and the translations between them are straightforward. In such regimes, a single manifold model is not only convenient; it is earned.

This reframing helps clarify what it means to say "GR explains gravity." It means: in the regimes where translations are integrable, GR is the compressed representation of the translation connection. It also clarifies what it would mean for GR to fail. GR would fail not because geometry is wrong everywhere, but because the translation structure ceases to be integrable in

certain regimes. In those regimes, insisting on a single manifold may be an overreach, and the correct global object may be twisted, layered, or context-indexed in a way that a single metric cannot capture without auxiliary witnesses.

In this sense, contextual gravity does not compete with GR on its home turf. It explains why GR is so powerful and why it might not be globally sufficient in cosmology without additional structure.

6.4 Where deviations should appear first: outskirts, mergers, voids, multi-probe seams
If contextual gravity is more than a re-description, it must make disciplined expectations about where non-integrability (and therefore apparent "dark" residuals) should appear most strongly. The central heuristic is: deviations concentrate where overlaps are real but translations are hardest.

Four classes of seams are especially natural.

First, outskirts and low-acceleration regimes. The outer regions of galaxies and clusters are where dynamical inferences are most sensitive to modeling assumptions (equilibrium, anisotropy, tracer bias) and where baryonic distributions become sparse and uncertain. They are also where small changes in translation rules can produce large apparent mass discrepancies. In a contextual view, this is a high-curvature region of inference-space: transport from baryonic mass models to dynamical predictions becomes path-dependent.

Second, mergers and non-equilibrium systems. Cluster collisions and other violent events stress the overlap conditions because different components behave differently: gas is collisional and shock-heated; galaxies are approximately collisionless; lensing integrates everything along the line of sight. These are precisely the environments where translation maps that assume equilibrium or smoothness are most likely to fail. In a contextual framework, mergers are not just "tests of dark matter"; they are tests of whether translations between probes commute under extreme conditions. Large holonomies are expected if the naive global-section category is too rigid.

Third, voids and under-dense environments. Voids are a complementary seam: low density, different screening behavior in many modified-gravity ideas, different tracer systematics, and a different balance of baryonic processes. If the obstruction is tied to environment-dependent translation curvature, voids should show a characteristic pattern: not necessarily "more discrepancy," but a differently structured discrepancy that depends on which translations are being exercised.

Fourth, multi-probe seams themselves. The most direct deviations should be visible not only in any single dataset, but in the relationships between datasets. The hallmark of translation curvature is that two routes from context A to context C disagree. That means the sharpest

observables are of the form “cross-probe consistency conditions” rather than single-probe fits. In practical terms, contextual gravity should predict structured residual relationships between lensing-inferred potentials and dynamics-inferred potentials as a function of system class, environment, and scale. The location of deviation is therefore not merely “at radius r,” but “at the interface between inference pipelines.”

This list also clarifies what would count as failure of the proposal. If one can define reasonable translation rules and find that loops commute to high accuracy across these seam classes without invoking witness objects, then the obstruction hypothesis collapses; the residuals would have to be explained as ordinary systematics or as ordinary hidden mass in the traditional sense. Conversely, if loop residuals persist and display stable patterns across systems and probes, then the curvature-of-translation picture gains empirical traction, and “dark matter” becomes increasingly well-described as an obstruction class rather than as a single ontological commitment.

This section therefore supplies the bridge between metaphor and program. Gravity, in this paper’s proposal, is the integrability structure of inference across contexts. Curvature is measured by path dependence of translations. GR is the integrable limit where translations can be compressed into a smooth manifold description. Deviations appear first where translations are hardest: outskirts, mergers, voids, and multi-probe seams.

7. Quantum Nonlocality Reframed: No Global Section Instead of “Spooky Influence”

The phrase “spooky action at a distance” captures a genuine shock: quantum theory predicts correlations that do not admit the kind of classical causal story our intuitions prefer. But “spooky influence” is not a diagnosis; it is a narrative placeholder. A topos-like approach offers a sharper diagnosis with less metaphysical drama: the problem is not that influences travel faster than light; the problem is that there is no single global assignment of classical properties that simultaneously agrees with all measurement contexts. In other words, the structure fails to glue.

This reframing aligns quantum foundations with the paper’s broader thesis. Cosmology’s dark-sector puzzles, in this view, may not be isolated anomalies but echoes of a general phenomenon: locally coherent descriptions that resist global sectionhood when we insist on too rigid a category of “one true underlying thing.” Quantum mechanics already taught physics that “what is real” is not always representable as a global valuation over all contexts. The proposal here is not to claim that dark matter is quantum contextuality; it is to claim that both may be instances of the same structural motif: obstruction to gluing under natural translation rules.

The point of this section is therefore twofold. First, it restates quantum nonlocality in the language of contextual sections and gluing impossibilities. Second, it uses that restatement as a conceptual bridge: if no-global-section is an acceptable, testable, and already-established feature in quantum theory, then "no global section" is no longer an exotic claim when invoked in cosmology. It becomes a plausible hypothesis about why global reconstructions can be brittle.

7.1 Measurement contexts as opens: outcomes as local sections

A measurement context is a specification of what is being measured and how: which observables are jointly measurable, which settings are chosen, and which experimental constraints define the sampling. In a topos-like perspective, contexts play the role of "opens" in a coverage: they are domains in which certain statements are meaningful and testable.

Within each measurement context, the theory provides probabilities for outcomes and, in some cases, functional relationships among compatible observables. These outcomes can be treated as local sections: assignments of values to the observables that are defined in that context. The crucial point is that a local section is not a full global assignment. It is an assignment only for that context's jointly measurable set.

The classical intuition expects that these local sections come from restricting a global section. That global section would be a hidden assignment of definite values to all observables, whether or not they are simultaneously measured, such that every context's observed outcome is merely the restriction of that global assignment to the measured subset. This is exactly the "realism" part of local realism, and it is also the implicit metaphysics behind the "spooky influence" narrative: if the values are real and global, then the weirdness must come from influences that coordinate them across space.

The topos-like view says: do not assume the global section exists. Treat each context as a legitimate local domain of truth, and ask whether the family of local sections can glue. The remarkable content of quantum theory is that in certain families of contexts, gluing is impossible. That impossibility is not a story; it is a mathematical constraint encoded in the structure of the theory and borne out by experiment.

7.2 Bell and Kochen–Specker as gluing impossibilities

Bell-type results and Kochen–Specker-type results are often presented as philosophical puzzles about reality. In the present framework, they are technical statements about gluing.

Bell-type results, in broad terms, show that no theory can reproduce all quantum predictions while satisfying both locality (no superluminal causal influence) and realism (pre-existing global values that determine outcomes) under standard assumptions about freedom of settings and statistical structure. In the gluing language, what fails is the existence of a global section that simultaneously restricts to all the relevant local sections while preserving the locality constraints

on how those sections may be correlated. The experimental violations of Bell inequalities are empirical evidence that the naive global-section model is not available.

Kochen–Specker results sharpen a different edge. They show that for certain quantum systems, it is impossible to assign noncontextual definite values to all observables in a way that preserves functional relations among commuting observables. In gluing language, even if one ignores spatial separation and focuses only on compatibility structure, there is still no global valuation that restricts consistently to all measurement contexts. The obstruction is internal: it is a feature of the lattice of contexts itself.

The key takeaway is that quantum theory's "weirdness" can be described without appealing to mysterious influence. The world permits coherent local sections in each context, but the family of contexts has no global section in the classical sense. "Nonlocality" becomes a misleading label if it suggests superluminal signals. The more accurate label is non-gluability: the failure of classical global valuation.

In the spirit of this paper, this is a decisive precedent. It demonstrates that physics already contains a successful theory in which "no global section" is not an admission of defeat but a precise, predictive structural fact.

7.3 Preserving causality without classical realism: signaling vs correlation

One reason the "spooky influence" narrative persists is that people conflate correlation with communication. Quantum entanglement produces correlations that cannot be explained by local hidden variables, but it does not thereby grant controllable faster-than-light signaling in standard quantum theory. This distinction matters because it clarifies what is actually sacrificed when one abandons the classical global section.

The sacrifice is classical realism in a particular form: the idea that there exists a context-independent global assignment of definite values to all observables that predetermines outcomes. Once that is relinquished, the need for superluminal influence evaporates. Outcomes are not the restriction of a global classical state of affairs; they are context-indexed events whose correlations are governed by the theory's structure.

Causality, in the operational sense relevant to relativity, is preserved because the rules do not allow an agent to use entanglement correlations to send a chosen message faster than light. The theory constrains what can be controlled. The correlations are real and experimentally robust, but they are not a communication channel. The basic lesson is: abandoning global classical realism is not abandoning causality; it is abandoning a specific classical story about how causality must look.

This lesson generalizes directly to cosmology. If a model's global section fails to exist, one is not forced into nihilism. One is forced into a more careful distinction between what is locally

operational, what is transported, and what is globally assertable. Quantum theory teaches that this distinction can be both precise and predictive.

7.4 A bridge concept: contextuality and dark-sector obstructions as the same kind of failure
The bridge proposed here is structural, not literal. The claim is not that dark matter is quantum entanglement, or that galaxy rotation curves are explained by Kochen–Specker. The claim is that both domains may share a common mathematical pattern: locally coherent descriptions that resist being unified as restrictions of a single global object under the assumed translations.

In quantum theory, the local objects are measurement-context outcome assignments (or probability assignments), and the translation constraints encode compatibility and functional relations. The obstruction manifests as the impossibility of a global valuation. In cosmology, the local objects are probe-specific gravitational inferences (dynamics, lensing, gas, timing), and the translation constraints encode how these inferences should agree on overlaps if they arise from a single global gravitational structure. The obstruction manifests as a persistent residual—often labeled dark matter or dark energy—needed to force agreement under a naive global-section assumption.

Seen this way, "dark matter as gluing obstruction" is not an arbitrary metaphor imported from mathematics. It is a disciplined extension of a pattern already present in physics. Quantum theory tells us that insisting on global classical valuation is too strong; the world is contextual. Cosmology may be telling us that insisting on a single globally integrable gravitational description in our chosen category may also be too strong; the world may be contextually glued with nontrivial twists.

This bridge also provides a practical research heuristic. In quantum foundations, progress came from identifying exactly which family of contexts makes gluing impossible and from quantifying the obstruction via inequalities and theorems. The analog in cosmology would be to identify which family of observational contexts produces non-gluability and to quantify it via cross-probe commutativity tests, loop-holonomy residuals, and obstruction-class classification. The expected outcome is not metaphysical clarity; it is a sharper map of where the global-section demand fails, and therefore where new structure must be introduced.

Finally, the bridge reduces the psychological burden of adopting a topos-like stance. If "no global section" is already an accepted feature of the microphysical world, then it is not an exotic surrender to suggest that large-scale cosmological modeling may also require a weaker notion of globality. The intellectual move is then consistent across scales: treat reality as locally well-defined within contexts and demand explicit translation and gluing rather than assuming that a single global story must exist by default.

This section thus completes a key part of the paper's argument. The nonclassical features of quantum theory are best understood as gluing impossibilities rather than spooky influence. That precedent legitimizes the paper's central proposal: treat cosmological dark-sector residuals as potential gluing obstructions in a sheaf of gravitational inferences, and treat the search for "new physics" as the search for the minimal additional structure that restores commutativity—or, more radically, as the recognition that the correct global object is inherently twisted and cannot be reduced to a single section in the naive category of classical models.

8. Multi-Scale Topoi: Renormalization as Functorial Flow

If the Universe is treated as a family of rule-contexts, then "scale" is not a minor label; it is one of the primary axes along which contexts differ. The rules that are expressible, the observables that are stable, and the symmetries that are meaningful all depend on the resolution at which you interrogate the world. Physics already knows this in practice through effective field theory and renormalization, but it often retains an implicit metaphysical picture: there exists one ultimate set of laws, and the rest are approximations. A topos-like approach does not deny the possibility of an ultimate law, but it shifts the emphasis. It treats the scale hierarchy itself as an organizing structure: a tower of contexts indexed by scale, with explicit maps between them. The flow of descriptions across scale is then not an afterthought; it is the central dynamical operation that produces what we call "laws."

In this section, renormalization is recast as functorial flow: a disciplined way to transport objects and truths from fine contexts to coarse contexts while tracking which structures survive, which symmetries emerge, and which obstructions appear. This recasting has two strategic payoffs. First, it provides a formal home for emergent symmetry and approximate laws. Second, it supplies a natural mechanism for "dark sector" phenomena to arise not as new substances but as anomalies of commuting diagrams: extra terms required to make multi-scale translations consistent. Finally, it offers a conceptual route toward the GR/QM mismatch: the possibility that GR and quantum theory are not two rival descriptions in one logic, but two internal logics at different scales connected by structured functors.

8.1 A topos tower indexed by scale: coarse-graining as a functor

The basic object here is a scale-indexed family of contexts. Think of a parameter s representing resolution or energy scale. For each s, there is a context category C_s (or a topos-like semantics T_s) describing what can be meaningfully stated and measured at that scale. The collection {T_s} forms a tower.

The essential ingredient is not the tower alone but the maps between levels. Coarse-graining is not merely throwing away information. It is a structured operation that takes a fine description

and produces a coarse one: it maps fine observables to coarse observables, fine states to coarse effective states, and fine dynamics to coarse effective dynamics. In categorical language, coarse-graining is a functor from the fine context to the coarse context. It transports not only objects but also the relations among objects.

This makes scale transition explicit and testable. In ordinary physics, we often speak as if "the same system" is being described at multiple scales. But what is the same is not the raw description; what is the same is the existence of a functorial relationship that preserves certain invariants. The tower view says: the identity of "the system" across scales is the coherence of the coarse-graining functors.

Once coarse-graining is functorial, two hazards become visible.

First, not all coarse-grainings commute. If you coarse-grain in one order (first average over microstates, then integrate out fast modes) you may not end up with the same effective description as if you do it in another order. That path dependence is the multi-scale analog of curvature in translation rules. The tower view therefore predicts that "effective laws" can depend on how you traverse the scale landscape, unless integrability conditions are satisfied.

Second, some structures exist only at some levels. Certain observables are definable only at coarse scale; certain micro-details are not stable under coarse-graining. The tower view makes this normal. A claim is not "globally true" in an absolute sense; it is true in T_s and possibly transported to other levels by functors. The question is what survives transport.

8.2 Universality as stable structure under functorial descent

One of the deepest successes of modern physics is universality: radically different microscopic systems can exhibit the same macroscopic behavior. In renormalization language, this happens because many micro-details wash out under coarse-graining and flows approach fixed points or stable manifolds in theory space. In a multi-scale topos view, universality is the statement that certain structures are stable under functorial descent.

This gives a clean, non-metaphysical definition of "law-like." A "law" is not necessarily a microscopic truth; it is a structure that persists under transport across relevant contexts, including across scale. If many different fine contexts map into the same coarse structure via coarse-graining functors, then the coarse structure is universal. It becomes a robust object precisely because it is insensitive to micro-variation.

This reframing has two implications that matter for cosmology.

First, symmetry becomes emergent in the strongest sense. A symmetry may not exist at fine scale, but if coarse-graining functors systematically erase the symmetry-breaking details, then symmetry appears at large scale as a stable structure under descent. This is exactly how many

macroscopic symmetries behave: they are not imposed; they are what remains after information loss.

Second, universality is what grants legitimacy to global cosmological parameters. A parameter is meaningful not because it is metaphysically fundamental, but because it remains stable under variations in micro-structure and under changes in context. The tower view therefore turns "what we can measure" into "what survives descent," which is a precise and useful criterion.

8.3 Dark sector as an anomaly of commuting diagrams: the "counterterm" analogy

The most provocative payoff of the multi-scale view is a new way to interpret the dark sector. In effective field theory, when you integrate out degrees of freedom, you often generate new terms in the effective action. These terms are not arbitrary; they are required to preserve consistency, renormalizability, or symmetry under the coarse description. In perturbative calculations, you sometimes need counterterms to cancel divergences and make predictions finite and consistent. The counterterms are not "new substances"; they are consistency restorers in the effective description.

A multi-scale topos suggests an analogous possibility for cosmology. Suppose you have a diagram of translations across contexts and scales. There are multiple ways to get from a fine description to a coarse observational prediction: you can coarse-grain first, then translate to an observational context; or you can translate first, then coarse-grain; or you can traverse a different chain of intermediate contexts. If the diagram commutes, you get consistent predictions. If it does not commute, you get residuals.

In the present proposal, the dark sector can be interpreted as an anomaly term introduced to restore commutativity. "Dark matter" and "dark energy," in this view, are not necessarily physical substances added to the Universe; they may be the effective terms required to make multi-scale translation diagrams commute under the modeling category we have chosen. They function like counterterms: additions that ensure the effective, coarse-grained description can be glued consistently across contexts.

This analogy is not meant as a rhetorical flourish; it suggests concrete expectations:

- The "dark" term should correlate with the structure of non-commuting diagrams, not simply with local mass density.
- The inferred dark component may be probe-dependent and scale-dependent in a systematic way, reflecting the specific path dependence of the translation chain.
- The term should be more prominent where coarse-graining is most aggressive or where intermediate contexts introduce large degeneracies.

- In regimes where the diagrams commute naturally (for example, well-tested local gravity contexts), the dark term should vanish or become tightly constrained.

This transforms the dark sector from a metaphysical embarrassment into a diagnostic tool. You do not begin by asking "what is the particle?" You begin by asking "which diagrams fail to commute, by how much, and what minimal completion makes them commute?" If the minimal completion looks like a stable, transferable object across contexts, it may deserve ontological status. If it looks like a context-dependent patch, it remains an effective counterterm.

8.4 GR/QM mismatch as incompatible internal logics connected by scale functors

The final move of this section addresses the most stubborn seam: the mismatch between general relativity and quantum theory. In the standard story, GR and quantum theory are two theories about the same reality, and therefore they must ultimately unify into one framework expressed in one logic. The multi-scale topos view allows a different posture: GR and quantum theory may be descriptions that live naturally in different levels of the tower, with different internal logics, connected by scale functors rather than merged by brute force into a single global formalism.

In quantum theory, the internal logic is contextual: not all propositions have global truth-values independent of measurement context. In GR, the logic is closer to classical field realism: fields and geometry are treated as globally defined structures on a manifold. The mismatch is often framed as technical, but it is also logical. A topos-like program makes that explicit: the two theories may be using different internal logics because they are native to different context strata.

If that is true, then the unification problem changes shape. Instead of demanding one global theory that reduces to both, one demands coherent functorial connections between the strata: maps that carry quantum contextual descriptions into classical geometric descriptions under coarse-graining and decoherence, and maps that carry classical geometric constraints back into the quantum domain as boundary conditions, selection rules, or consistency constraints. The requirement becomes: the scale functors must preserve the right invariants and must make the relevant diagrams commute in the regimes where both descriptions apply.

This is not a soft option. It is a demanding one, because it insists that the emergence of classical spacetime geometry be described as a functorial descent from a more contextual logic, and that the return influence of geometry on quantum processes be described as a disciplined pushforward or pullback of constraints. But it has a conceptual advantage: it allows one to accept that "one logic to rule them all" may be a psychological projection, while still demanding that the relationship between logics be precise.

In this frame, the GR/QM mismatch is not merely a missing equation. It is a missing account of how internal logics relate across scale. If the Universe is an informational dynamic topos of rules, then “quantum” and “geometric” may be two stable regimes of rule-logic in the tower, and the deep problem is to write down the functors that connect them.

This completes the multi-scale synthesis. Renormalization becomes functorial flow in a tower of contexts. Universality becomes stability under descent. The dark sector becomes a candidate anomaly term required to make multi-context, multi-scale diagrams commute. And the GR/QM mismatch becomes a mismatch of internal logics whose resolution may be a network of scale functors rather than a single monolithic theory written in a single presumed logic.

9. Falsifiability: What This Framework Must Predict (and Risk)

A topos-like cosmology only becomes physics when it accepts the same bargain as every serious theory: it must compress the world with fewer knobs than its competitors, and it must risk decisive failure. The obstruction picture makes a strong claim that can be tested: the dominant “dark” discrepancies arise because local gravitational inferences from different contexts do not glue into a single global section under the assumed translation rules, and the mismatch has structured, repeatable form. That claim is not automatically true. It will fail if improved modeling and calibration make cross-probe translations commute without introducing new gluing structure, or if a particle-like completion explains the same cross-probe patterns with fewer degrees of freedom and higher predictive stability. This section defines the minimum set of predictions and risks that “obstruction cosmology” must take on to be meaningfully falsifiable.

9.1 Cross-probe residual structure: mandatory relations between lensing and dynamics

The framework’s most direct test is not “does it fit rotation curves?” or “does it fit lensing maps?” in isolation. Its core claim is about overlaps, so it must predict relations between probes on the same systems.

In ordinary language, dynamics measures how matter moves and lensing measures how light bends. In many standard formalisms these correspond to different combinations of underlying gravitational potentials (often discussed as a potential that governs time dilation and a potential that governs spatial curvature). If a single global gravitational structure exists and the translations are integrable, then the two probes are not free to disagree arbitrarily. Their relationship is constrained.

Obstruction cosmology therefore owes “mandatory residual structure” of the following kind:

First, residuals should be diagrammatic, not ad hoc. If you define two translation routes from a dynamical inference to a lensing prediction (for example, dynamics -> inferred potential ->

predicted lensing, versus dynamics -> baryon-calibrated mass model -> predicted lensing), the difference between routes should follow a stable rule. The residual should look like a holonomy, not like a per-object patch.

Second, residuals should be probe-covariant. If you change lensing reconstruction method within a defined equivalence class, or you change dynamical tracer selection within a defined equivalence class, the obstruction should transform predictably rather than changing arbitrarily. If the "dark" component is merely a proxy for uncontrolled systematics, it will not have this covariance.

Third, residuals should have constrained functional form across systems. The framework must forbid a free function per galaxy. It should instead predict that the lensing-dynamics mismatch depends on a small set of context features such as scale regime, equilibrium stress, environment density, and coarse-graining level, with a shared parameterization reused across many systems.

Fourth, residuals should respect calibrated null regimes. In regimes where translation commutativity is already strongly supported (high-acceleration, well-characterized baryon distributions, and clean lensing geometries), the framework must predict small or vanishing mismatch. If it predicts a large "obstruction" everywhere, it becomes unfalsifiable and loses explanatory power.

In short, the falsifiable content is not "there is a mismatch," but "the mismatch has a constrained, reusable structure linking lensing and dynamics across many systems and context changes."

9.2 Cluster collisions as separation tests: where obstructions must localize

Cluster mergers and collisions are valuable because they bring multiple contexts into hard contact under conditions that break simplifying assumptions. They are where "gluing" is most likely to fail if gluing is nontrivial.

The obstruction framework must make clear, risky commitments about localization:

First, localization in overlap, not in narrative. The framework predicts that obstructions concentrate where translation maps are stressed: where equilibrium assumptions fail, where projection effects dominate, where collisional and collisionless components decouple, and where line-of-sight integration confounds naive identification of "the same mass." The obstruction should therefore correlate with measurable stress indicators (e.g., dynamical disequilibrium proxies) rather than merely with radius or total mass.

Second, localization in diagrams. A cluster collision is a multi-route test. You can infer gravitational structure via lensing, via galaxy dynamics, and via gas modeling. Obstruction

cosmology predicts that the largest residuals will appear specifically on loops that traverse contexts with incompatible assumptions (for example, loops that rely on hydrostatic equilibrium during violent merger phases). The mismatch should not be equally large in all routes.

Third, localization that is reproducible across event classes. If “dark matter” is an obstruction class, different merger geometries should produce different holonomy patterns, but those patterns should belong to a small taxonomy. A theory that cannot predict which merger classes should amplify which residual loops is not predictive enough.

The risk is sharp: if cluster collisions show that all well-defined translation loops commute once systematics are properly modeled and calibration is consistent, then the obstruction interpretation becomes unnecessary. Conversely, if residual holonomies persist and are structured by loop choice and event class, obstruction cosmology gains a unique, falsifiable foothold.

9.3 Null predictions: when the obstruction must vanish (and why that matters)

Null predictions are where this framework either becomes scientific or collapses into unbounded flexibility. It must specify conditions under which the obstruction is predicted to be zero or negligibly small.

Several null regimes are natural and should be explicit:

First, high-control local regimes. In contexts where clocks, rods, and energies are transported with exceptional calibration control and where gravitational dynamics is tested at high precision, the translations should be integrable to within known tolerances. The obstruction must be constrained to be tiny there. This is not merely respecting existing successes; it is the framework’s own internal requirement that curvature of translation rules is not arbitrarily large.

Second, single-context closures. If a system is probed in a way that does not require cross-context translation (or requires only a well-validated translation), the framework should not manufacture an obstruction. This creates a clear diagnostic: obstruction arises from overlap complexity, not from the mere presence of gravitation.

Third, symmetry-stable regimes under specified coarse-grainings. If the framework claims certain symmetries are emergent at large scale under particular coarse-graining functors, then within those coarse-grained contexts the obstruction should vanish in the corresponding symmetry channels. This gives testable “channel nulls”: certain kinds of anisotropy or slip should be suppressed if the emergent symmetry claim is correct.

Fourth, commuting-subdiagram nulls. Even if the full network of contexts does not glue globally, there should be subfamilies of contexts for which gluing does hold. The framework must

identify these and predict that the obstruction class restricted to these subcovers is trivial. This is essential: without commuting islands, the theory becomes an all-purpose excuse generator.

Nulls matter because they prevent the framework from becoming a relabeling of ignorance. A good obstruction theory says not only "here is where gluing fails," but also "here is where gluing must succeed, and by how much."

9.4 Model selection penalties: degrees of freedom and the "no free function per object" rule
The dominant failure mode of alternatives to standard cosmology is not incorrectness but non-constraint: they can explain anything by letting something vary arbitrarily. Obstruction cosmology must therefore impose a strict complexity discipline.

The core rule is simple: no free function per object. Any additional flexibility introduced to restore gluing must be shared across objects by functorial reuse. Concretely:

First, the obstruction must be parameterized by a small set of global parameters or by a tightly regularized function class that is common across systems. If each galaxy or cluster gets its own bespoke correction, the theory has no predictive content.

Second, model comparison must be carried out on out-of-sample cross-probe prediction, not in-sample fit. A theory should be rewarded for predicting lensing from dynamics (or dynamics from lensing) on new systems with fixed global settings. This is where obstruction cosmology claims its advantage: it targets the translation layer, so it should excel at cross-context prediction if it is right.

Third, penalties must be applied where they belong: on morphisms and gluing structure. If the theory modifies translation rules between contexts, that modification is counted as complexity. It cannot be hidden as "just a calibration choice." The framework's honesty depends on this.

Fourth, uncertainty transport is part of complexity. If the theory only "works" by inflating uncertainties in one probe or by selectively down-weighting overlaps, that must be counted as loss of explanatory power, not as success.

This discipline is itself falsifiable. If a particle dark matter model or a well-posed modified gravity model achieves stronger cross-probe predictive performance with fewer degrees of freedom, obstruction cosmology loses.

9.5 Distinguishing three programs: particle dark matter, modified gravity, obstruction cosmology
To be scientifically useful, the framework must not merely critique; it must help distinguish competing explanations by prediction patterns. The three major programs differ most sharply in what they treat as primary and where they locate "new structure."

Particle dark matter program
Primary claim: there is additional matter content with gravitational influence that couples in a mostly universal way, independent of probe choice.
Typical signature: a probe-independent mass distribution that simultaneously reconciles dynamics and lensing when baryonic physics is modeled, with residuals shrinking as data improve.
Key risk: if cross-probe mismatches show structured probe-dependence that cannot be captured without adding extra non-gravitational couplings or context-dependent behavior, the program becomes more complex than advertised.

Modified gravity program
Primary claim: the gravitational law or field equations differ from the standard form, often in a regime-dependent way (acceleration scale, environment screening, etc.).
Typical signature: a predictable relationship between baryonic distribution and gravitational response that reduces the need for unseen matter, often with characteristic environment dependence.
Key risk: if the required modification effectively becomes a system-specific function, predictive content collapses, and cross-probe consistency can break in uncontrolled ways.

Obstruction cosmology program
Primary claim: the dominant “dark” residuals arise because local gravitational inferences do not glue into a single global section under naive translation rules; the mismatch is a structured obstruction concentrated at overlap seams and along specific translation loops.
Typical signature: persistent, repeatable, loop-dependent holonomies linking lensing and dynamics, with a small taxonomy of obstruction patterns across system classes, plus explicit null regimes where commutativity must hold.
Key risk: if improved modeling makes loops commute broadly without extra gluing structure, or if the required obstruction becomes per-object tuning, the program fails.

A decisive discriminator, in this framing, is the topology of residuals: do residuals look like “missing mass” (substance), like “law deformation” (modified dynamics), or like “non-commuting translations” (obstruction/holonomy)? Obstruction cosmology is committed to the third. If the residuals do not organize into loop-structured, reusable patterns, the proposal is not merely weakened, it is refuted as a distinct explanatory program.

This section therefore sets the standard the remainder of the paper must satisfy: identify the loops, define the overlaps, compute the residual structures, publish null predictions, and accept model-selection penalties. Anything less is philosophy.

10. A Minimal Research Program and Prototype Implementation

A proposal like "dark matter as a gluing obstruction in a topos-like family of contexts" earns its keep only if it can be implemented with today's data and judged by clear outputs. The goal of a minimal research program is not to prove the metaphysics. It is to build a prototype that takes real astronomical datasets, constructs explicit context objects, defines explicit morphisms, and then computes whether the resulting diagrams commute. The deliverable is not a new narrative; it is a set of obstruction diagnostics: loop residuals, localization maps, and null-regime confirmations or violations.

This section specifies a practical starting point. It deliberately avoids maximal generality. The aim is to get to a working pipeline that can be shared, rerun, and falsified by other groups. If the framework has substance, it should show up as stable, reusable patterns in the computed obstructions across systems and probes. If it does not, the program should fail quickly and cleanly.

10.1 Define the context category: a practical first choice for astronomy

The first implementation choice is the context category: what counts as a "context," and what counts as a morphism between contexts. If this is defined too abstractly, the project becomes untestable. If it is defined too narrowly, it becomes a restatement of existing practice. The minimal choice should balance both.

A practical first context category for astronomy can be defined by three axes:

First, probe modality. At minimum, include:

- Dynamics context (D): rotation curves, stellar velocity dispersions, satellite kinematics
- Weak lensing context (WL): shear and convergence maps, stacked lensing profiles
- Strong lensing / time delay context (SL): when available for select systems
- Gas context (G): X-ray / SZ pressure profiles with equilibrium modeling for clusters

Second, system class. Begin with a small number of system types to avoid uncontrolled heterogeneity:

- Spiral galaxies with high-quality rotation curves
- Galaxy-galaxy lensing samples aligned with those spirals (or statistically matched populations)
- Clusters with lensing + X-ray/SZ + galaxy dynamics

Third, modeling regime tags. Each context object must carry a tag bundle that specifies the assumptions under which its inference is valid, for example:

- Equilibrium assumption strength (e.g., “approximate circular equilibrium” or “merger/nonequilibrium”)
- Baryon modeling family used (e.g., stellar M/L priors, gas mass models)
- Reconstruction method family (e.g., shear calibration pipeline class)

With these axes, a context becomes something like: “spiral galaxy dynamics under equilibrium modeling family A with baryon prior family B,” or “weak lensing stacked shear for matched spirals under shear pipeline family C.” Contexts overlap when they refer to the same physical systems or statistically matched ensembles and when there exists a declared translation that maps their outputs into a common representational space.

Morphisms in this category are of two main kinds:

- Probe translation morphisms: maps from a gravitational structure inferred in one probe to the predicted observables in another probe (or to a common latent structure representation).
- Coarse-graining morphisms: maps between individual-system contexts and ensemble/stacked contexts, or between high-resolution and smoothed descriptions.

This definition is concrete enough to implement: it mirrors how observational cosmology already segments analyses, but it makes the segmentation explicit and therefore testable.

10.2 Data objects and morphisms: what to store, what to map, what to compare

The prototype requires disciplined data objects. Each context must output more than a best-fit curve. It must output a structured object that can be transported and compared.

A minimal context object should include:

First, a latent gravitational representation. Choose a representation that can host both dynamics and lensing inferences. For an initial prototype, use one of these two options:

Option A: Effective radial acceleration or potential gradient profile g(r) with uncertainty and covariance, plus a map from g(r) to each probe’s observables under declared assumptions.
Option B: A pair of potentials (Phi, Psi) or an equivalent “lensing vs dynamical potential” decomposition, represented in a basis (e.g., radial bins or low-order basis functions) with covariance.

The specific choice matters less than consistency. The goal is to have a common target space for morphisms.

Second, a forward model. For each context, store the forward model that maps the latent representation to the raw observables of that probe. This includes nuisance parameters and priors that are explicitly declared.

Third, an inference output as a distribution, not a point. Store posterior samples or a Gaussian approximation (mean + covariance) in the latent representation space. Without uncertainty structure, commutativity tests are meaningless.

Fourth, provenance metadata. Store the modeling family tags, calibration references, selection criteria, and any known systematics flags. In a topos-like framework, these are not “notes”; they are part of the context definition.

Now the morphisms.

A morphism is a map that takes one context’s latent object into another context’s latent object or into the other context’s observable space. The minimal set includes:

- Dynamics-to-latent morphism: observed velocities -> inferred g(r) or (Phi, Psi) under equilibrium assumptions and baryon model.
- Lensing-to-latent morphism: observed shear/convergence -> inferred lensing potential or projected mass proxy -> latent representation.
- Gas-to-latent morphism: pressure/temperature profiles -> inferred potential under equilibrium model -> latent representation.
- Coarse-graining morphism: individual latent profiles -> stacked latent profile with correct weighting and covariance propagation.

What to compare.

Comparisons should happen on overlaps, and overlaps must be defined operationally:

- Same system overlaps: a specific galaxy or cluster with multiple probes.
- Matched ensemble overlaps: statistically matched samples where individual overlap is not possible, but ensemble overlap is meaningful.

In each overlap, comparisons should be computed in the same latent space after applying the morphisms.

10.3 Computing obstructions from real datasets: a step-by-step pipeline

The prototype obstruction computation can be organized as a loop test. The output is a holonomy: the residual after transporting around a loop of contexts.

A minimal step-by-step pipeline:

Step 1: Select a small, clean dataset bundle
Pick:

- A set of spiral galaxies with high-quality rotation curves and baryonic maps
- A weak-lensing dataset that can provide shear profiles for a matched sample (stacked if necessary)
- A set of clusters with lensing + X-ray/SZ + galaxy velocity dispersion data

Keep the first pass small. The goal is not statistical dominance; it is structural diagnosis.

Step 2: Define the latent space
Choose g(r) in radial bins (or a low-order basis) as the latent representation. For clusters, allow projected quantities if required, but always map into a shared representational target for comparison.

Step 3: Build context objects
For each probe and system class:

- Run the inference to obtain posterior samples for the latent representation
- Record forward model family and nuisance parameters
- Store posterior mean/covariance (or samples)

Step 4: Specify morphisms
Explicitly write the maps between contexts, including any approximations:

- Map dynamics latent g(r) to lensing predictions and invert if necessary (or map both to a shared potential basis)
- Map gas-inferred potentials to the same latent basis
- Map stacked lensing results to an effective latent g(r) for the ensemble

Step 5: Define loops
A loop is a cycle of morphisms that begins and ends in the same context space. Example loops:

Loop A (galaxy-scale, ensemble):
Dynamics (spirals) -> latent g(r) -> predicted lensing (for matched sample) -> inferred lensing latent g(r) -> compare back to dynamics latent

Loop B (cluster-scale):
Gas -> latent -> predicted lensing -> inferred lensing latent -> predicted dynamics -> inferred dynamics latent -> compare back to gas latent

Loop C (method robustness loop):
Lensing pipeline A -> latent -> lensing pipeline B -> latent -> compare (tests morphism equivalence within lensing contexts)

Step 6: Compute holonomy residuals
For each loop, compute the residual between the starting latent object and the returned latent object. This residual is the obstruction witness for that loop. Quantify it with:

- A normed residual statistic using full covariance
- A radial or spatial localization map: where in the profile does the residual concentrate?
- A dependence on context tags: does the residual correlate with disequilibrium indicators, environment, scale?

Step 7: Test null regimes
Before claiming "obstruction," identify subcovers where translations are expected to commute (high-quality, equilibrium-like systems; regimes where modeling is strongest). Confirm that holonomies are near zero there. If not, you likely have a pipeline or modeling problem, not new structure.

Step 8: Cross-validate across systems
The framework only gains credibility if the residual structure repeats across systems in a small taxonomy. If every system yields a unique residual shape, you have fallen into the "free function per object" trap.

Step 9: Compare with competitor models
For the same dataset bundle, run a conventional particle dark matter model fit and a representative modified gravity fit (with disciplined complexity control). Compare predictive performance specifically on cross-probe overlap tests, not merely on single-probe fits.

The output of this pipeline is not a philosophical conclusion. It is a set of empirical artifacts: residual maps, loop statistics, and a classification attempt.

10.4 Benchmarks and shared tasks: rotation curves + weak lensing + cluster maps
To make the program communal and falsifiable, it should be framed as shared tasks with clear inputs and outputs.

A minimal benchmark suite could include three tasks:

Task 1: Galaxy rotation curve to lensing cross-prediction
Input: rotation curves + baryonic mass models for a defined sample, plus lensing shear profiles for a matched sample (stacked if needed).
Output: predicted shear profile from dynamics-inferred latent structure, compared to observed

shear.
Score: cross-probe residual statistic and holonomy loop measure.

Task 2: Weak lensing to dynamics cross-prediction
Input: lensing profiles + baryon priors for matched sample, plus a subset of kinematic constraints.
Output: predicted rotation curve behavior or acceleration profile.
Score: same.

Task 3: Cluster multi-probe commutativity maps
Input: cluster lensing convergence maps, X-ray/SZ pressure profiles, and galaxy kinematics.
Output: spatial/radial map of loop residuals, plus correlation of residual hotspots with disequilibrium indicators.
Score: loop residual norms + localization agreement across event classes (relaxed vs merging clusters).

Each task should include baseline solutions:

- Particle dark matter baseline: a standard halo model with baryon modeling.
- Modified gravity baseline: a disciplined, low-parameter modification model.
- Obstruction baseline: one or two parameterizations of translation curvature or gluing obstruction, shared across systems.

The benchmark must reward out-of-sample transfer. A method that learns a bespoke correction per object should score poorly. A method that uses a small reusable structure to reduce cross-probe residuals across multiple datasets should score well.

10.5 Failure cases and refinement: when the framework collapses or becomes trivial
A serious research program must specify how it can fail, and what kinds of failure mean.

There are three main collapse modes.

Collapse mode 1: Everything commutes after careful systematics handling
If improved calibration, baryon modeling, and selection control make cross-probe loops commute broadly without introducing new structure, then “dark matter as obstruction” becomes unnecessary as a distinct explanatory framework. The residuals would be better interpreted as ordinary hidden mass in the standard sense or as resolved systematics. In this case, the program succeeds scientifically by being falsified cleanly.

Collapse mode 2: Residuals exist but are unstructured and non-transferable
If loop residuals persist but do not cluster into a small taxonomy and instead vary idiosyncratically by object, then the obstruction picture cannot be distinguished from ad hoc

modified models. It collapses into a per-object patch language, which violates the “no free function per object” rule. The correct conclusion would be that the framework has no predictive compression and should be rejected or radically constrained.

Collapse mode 3: Residuals are dominated by known modeling degeneracies
If the computed “obstruction” tracks known nuisance directions (e.g., mass-anisotropy degeneracy, shear calibration uncertainty, hydrostatic bias) and disappears when those are handled, then the obstruction is not new structure but a reflection of old uncertainties. This is not a failure of science; it is a success of diagnosis. It means the framework’s added value is epistemic bookkeeping rather than new physics.

Refinement pathways exist if the program does not collapse but also does not trivially validate:

- Improve context definitions: sharpen what counts as overlap and what morphisms are legitimate.
- Expand the global object class: allow twisted sections or stack-like completions if residuals are structured but cannot be removed by simple modifications.
- Tighten complexity penalties: force reuse and out-of-sample performance to prevent accidental overfitting.
- Identify decisive subcovers: find small sets of contexts where the obstruction signal is maximized and systematics are minimized.

The minimal research program is therefore deliberately austere. It translates the topos-like proposal into a computational pipeline that either produces transferable obstruction patterns or fails. That is exactly what a new framework should do: accept the risk of becoming unnecessary, while offering a clear, implementable path to demonstrating whether “dark matter” is better understood as an entity, a law modification, or a structured failure of gluing in the translation rules that connect our contexts.

11. Implications and Limits: What Changes, What Stays

A topos-like, descent-first cosmology is not a declaration that the past century of physics was misguided. It is a proposal to relocate where the discipline lives. Instead of treating global structure, symmetry, and self-consistency as foundations that must be assumed in order to do science, we treat them as achievements that must be earned by explicit gluing across contexts. That relocation changes how we interpret dark-sector residuals, how we think about “one Universe model,” and how we handle the seam between GR and quantum theory. But it does not abolish the successes of existing theories; it aims to explain why they work where they work, and why they may become epistemically brittle where they do not.

This section clarifies what must change if the framework is taken seriously, what can remain unchanged, and what the framework explicitly refuses to claim. The goal is to end with a realistic picture of what "success" would look like: not omniscience, but better compression and stronger cross-context prediction with explicit, testable limits.

11.1 The cosmological principle as a convenience prior, not a foundation

The cosmological principle—roughly, large-scale homogeneity and isotropy—has been one of the great simplifiers of cosmology. It makes the inverse problem solvable. It permits a small-parameter background model. It anchors the use of statistical summaries. None of that is denied.

What changes is the logical status. In a descent-first framework, the cosmological principle is treated as a convenience prior: a regularizer that we adopt because it is a powerful compression and because it appears to work well in certain contexts. But it is not a metaphysical foundation. It is not a "property the Universe must have so that the math makes sense." It is a hypothesis about which symmetries remain stable under the coarse-graining functors relevant to our observational context.

This change has practical consequences.

First, it makes the principle testable in the right way. Instead of asking whether the Universe is homogeneous "in itself," we ask: which coarse-grained observables are consistent with homogeneity and isotropy across the contexts we can actually compare? Where do those symmetries break, by how much, and in which channels? This is a better question because it respects the observational fact that we reconstruct from a thin sample of the past light cone.

Second, it blocks overreach. When we adopt the cosmological principle as a foundation, we tend to treat deviations as nuisances or as evidence that some additional component must be added to restore the global story. When we treat it as a convenience prior, deviations become information: they are potential evidence of nontrivial gluing, scale-dependent symmetry, or translation curvature. The principle remains valuable, but it is demoted from "truth" to "tool."

Third, it changes what counts as anomaly. A tension that disappears when you relax a convenience prior is not necessarily a failure of nature; it may be a failure of an enforced regularizer. Conversely, a symmetry that remains stable even when its breaking is allowed and penalized becomes stronger evidence for genuine universality.

So the cosmological principle stays in the toolkit, but it is no longer the unspoken metaphysical gatekeeper. It becomes one of the hypotheses that descent can certify or reject within declared contexts.

11.2 Realism without omniscience: global claims as derived, not assumed

A common fear when people hear “context-indexed truth” or “no global section” is that realism is being abandoned. The framework does not require that. What it requires is realism without omniscience: the refusal to equate “what is real” with “what admits a single global classical description available to our inference.”

In this posture, local reality is taken seriously. Outcomes happen. Signals arrive. Instruments respond. Regularities are discovered. But global claims—claims about the Universe’s full structure, its exact initial conditions, or its unique global geometry—are treated as derived objects. They are the result of successful gluing across contexts, not assumptions made at the beginning.

This flips the epistemic arrow. Instead of:

Assume a global model -> derive local predictions -> fit data -> interpret residuals as new stuff

we do:

Construct local models -> define translations -> test commutativity -> decide whether a global model is warranted -> interpret failures as obstructions with multiple completion options

In this view, a “single Universe model” is not the default; it is a certificate. When you can glue, you earn the right to speak globally. When you cannot glue, you do not deny reality; you admit that your access is structured and your global claims are not licensed.

This also reframes debates about “the Big Bang” as a model. The framework does not require rejecting such models, but it does insist that their global extrapolations be treated as derived compressions whose legitimacy depends on gluing stability across the observational contexts available. It is not anti-model; it is anti-unearned-globality.

11.3 Why this does not “solve everything”: humility about ontology and ultimate causes

A serious limitation of the approach is also its virtue: it is cautious about ontology. It does not promise to reveal what “dark matter really is,” what “gravity ultimately is,” or what “the Universe is in itself.” Those are ontological questions, and the framework is designed to keep them downstream of empirical structure rather than upstream as assumptions.

There are several reasons this humility is not evasion.

First, the obstruction framing separates phenomenology from completion. It prioritizes measuring the obstruction structure—where diagrams fail to commute, how residuals localize, what null regimes exist—before choosing a metaphysical story. This prevents premature reification of bookkeeping terms into substances, but it also means the framework may end with multiple plausible completions rather than one.

Second, it recognizes the inverse-problem limit. Even with perfect mathematics, finite data restricted to a thin observational surface can underdetermine global structure. A topos-like posture accepts that some global questions may remain underdetermined without new kinds of observational access. This is not pessimism; it is a recognition of what the measurement geometry permits.

Third, it does not claim to unify GR and quantum theory by decree. It offers a way to treat their mismatch as possibly involving different internal logics connected by scale functors, but it does not magically produce the missing functors. The framework is a research direction, not a final theory.

Fourth, it does not abolish the possibility that particle dark matter is real. It merely insists that the path to that conclusion should be via minimal completion of measured obstruction structure, not via reflex. If the minimal completion is particle-like and probe-independent, the framework will say so. If the minimal completion is translation-curvature-like, it will say that instead. The framework is not committed to one ontology; it is committed to the discipline of earning ontology from gluing success.

So the approach is not a universal solvent. It is an epistemic architecture designed to keep physics honest about what is assumed, what is transported, and what is earned.

11.4 What success would look like: fewer knobs, more cross-context predictive power

The success criteria are concrete. The framework succeeds only if it improves compression and prediction relative to competitors under strict complexity control.

Success would have at least four signatures.

First, a small taxonomy of obstruction patterns. Instead of an unstructured zoo of “dark profiles,” the computed holonomy residuals would cluster into a small number of repeatable classes associated with system type, environment, disequilibrium indicators, and scale. This would be the analog of universality classes in renormalization: stable patterns under functorial descent.

Second, cross-context prediction that beats ad hoc fits. The framework should predict lensing from dynamics (and vice versa) across new systems with fixed global settings and minimal tuning. It should reduce multi-probe tensions by modifying translation rules in a way that is shared across systems, not by adding bespoke per-object corrections.

Third, strong null regimes. The framework should specify and then confirm regimes where commutativity must hold. If it cannot make and pass null predictions, it is not falsifiable and not scientific.

Fourth, fewer knobs than the alternatives for the same predictive reach. Particle dark matter models often add new entities but can be parsimonious in parameters at the phenomenological level. Modified gravity models can be parsimonious but risk becoming free functions. Obstruction cosmology must show it can remain parsimonious by localizing complexity in a small number of translation-curvature or obstruction parameters that are reused across contexts and scales. If it requires more knobs than both competitors to match the same data, it fails.

If these signatures are met, the framework changes the culture of cosmology in a productive way. “Consistency” becomes descent, “symmetry” becomes commutativity under context change, “dark matter” becomes a measured obstruction rather than a default substance claim, and “global models” become derived compressions that must be earned. What stays is the scientific core: disciplined inference, transparent assumptions, and predictions that risk being wrong.

12. Conclusion: From Enforced Elegance to Earned Structure

Cosmology has achieved extraordinary predictive success with a remarkably compact set of ideas. That success is real and deserves to be preserved. But the very compactness that makes the field powerful also makes it epistemically fragile: when global conclusions are inferred from limited causal access, and when tractability depends on symmetry and consistency assumptions, it becomes easy to confuse enforced elegance with discovered structure. The proposal of this paper is a disciplined response to that fragility. It does not demand abandoning successful models. It demands relocating their authority. Instead of treating global structure as a starting axiom, it treats global structure as an earned certificate of gluing across contexts. Instead of treating symmetry as metaphysical invariance, it treats symmetry as commutativity under context-change maps. Instead of treating dark matter as an entity by default, it treats the dark sector first as an obstruction to descent. The conclusion is therefore not “throw away cosmology,” but “make cosmology honest about how it earns its global claims.”

12.1 Restating the pivot: models as glued rule-contexts

The pivot can be stated in one sentence: the Universe is approached as a family of rule-contexts, and “a model of the Universe” is the successful gluing of those contexts, not an assumed global story imposed at the beginning.

In standard practice, we begin with a global ansatz, then fit. In this proposal, we begin with local contexts—measurement modalities, scale regimes, causal access limitations, coarse-graining choices—and we build explicit morphisms that translate between them. Only if the translations

commute on overlaps do we declare the existence of a global section: a unified structure that restricts to each local description.

This reorganization changes what counts as fundamental. The fundamental objects become: (i) the contexts where claims are testable, (ii) the translation rules that move information between contexts, and (iii) the commutativity or failure of commutativity of the resulting diagrams. A global cosmological model becomes the *compression* of a network that actually glues. When it does not glue, the failure is not treated as embarrassment; it is treated as structured data.

In that sense, the proposal is an epistemic engineering project as much as a physical one. It aims to build a framework in which what is assumed is always labeled, what is transported is always explicit, and what is global is always earned.

12.2 Dark matter reinterpreted as obstruction, not entity

The dark matter problem is often narrated as a stark ontological choice: either there is unseen matter, or gravity is wrong. The obstruction framing adds a third, logically prior option: the mismatch is telling us that the local gravitational inferences from different contexts do not glue into a single global gravitational section under the translation rules we currently assume.

In this view, "dark matter" is first the name of a measured obstruction class: a persistent, structured failure of descent in the sheaf of gravitational inferences (dynamics, lensing, gas, timing). "Halos" become witnesses introduced to force gluing in a naive global-section category. Only after the obstruction is characterized do we ask what completion is minimal and stable: a particle-like entity, a modified law, a revised translation rule, or a broadened notion of global object (twisted sections, stack-like completions).

This does not rule out particle dark matter. It simply demotes it from default ontology to candidate completion. The scientific gain is a cleaner separation between what is empirical (non-gluability patterns) and what is interpretive (what the witness "really is"). That separation is exactly what prevents bookkeeping terms from hardening into metaphysical substances prematurely.

12.3 The new criterion: functorial predictability across contexts

If the framework is to matter, it must offer a criterion stricter than "fits the data." Many frameworks can fit the data when allowed enough flexibility. The proposed criterion is functorial predictability across contexts.

A theory earns credibility when:

- It defines how to transport inferences between contexts (morphisms).
- The transported predictions agree on overlaps (descent).

- The agreement persists under reasonable changes in context definition (naturalness).
- The complexity required to maintain commutativity is shared across systems (no bespoke per-object tuning).
- The theory makes strong null predictions: subcovers where obstruction must vanish.

This criterion directly addresses the psychological concern at the center of the paper. Symmetry and consistency are powerful regularizers. They are indispensable. But they must be treated as measurable outcomes of transport and gluing, not as metaphysical necessities. Functorial predictability is what distinguishes "enforced consistency" from "earned structure." It asks whether the same rule survives translation—across probes, across scales, across environments—without silently injecting priors at each step.

If this criterion is met, the framework strengthens realism rather than weakening it. It says: what is real is what remains stable under disciplined transport, not what we assume must exist globally because we prefer a single story.

12.4 Next steps: the smallest decisive tests

The path forward must be small, decisive, and falsifiable. The framework does not need a grand unification project to be judged. It needs a prototype that computes obstructions on real data and either finds transferable structure or fails.

The smallest decisive tests are therefore cross-probe commutativity tests with explicit loop diagnostics:

First, galaxy-scale loops linking dynamics and weak lensing
Build matched samples where rotation-curve-inferred latent gravity can be transported into predicted lensing profiles and compared against observed shear. The key output is not merely a fit but a loop holonomy: the residual after transporting around alternative routes. The decisive question is whether the residual has a stable, reusable structure across many systems under strict complexity control.

Second, cluster-scale multi-probe loops in relaxed vs merging systems
Compute loop residual maps combining lensing, gas modeling, and galaxy dynamics, and test whether residuals localize where translation assumptions are stressed (outskirts, mergers, disequilibrium zones) while vanishing in predicted null regimes. The decisive question is whether residuals behave like structured holonomies rather than like unstructured noise or per-object tuning.

Third, null-regime confirmations
Identify subcovers where commutativity should hold at high accuracy—high-control regimes with strong calibration and well-understood modeling—and confirm that the framework's

obstruction collapses there. If it does not, the framework is not diagnosing physics; it is diagnosing its own pipeline.

Fourth, model-selection head-to-head comparisons
Run the same cross-probe tasks under (i) a standard particle dark matter baseline, (ii) a disciplined modified gravity baseline, and (iii) an obstruction-cosmology baseline with strict penalties against per-object flexibility. Judge them by out-of-sample cross-context prediction, not by in-sample curve fitting.

These tests are “small” in the right way: they do not require new observatories, only disciplined bookkeeping and explicit translation rules. They are decisive because they can refute the framework: if loops commute broadly once systematics are handled, obstruction cosmology is unnecessary as a distinct program. If loops fail to commute but the residuals are non-transferable, the framework collapses into ad hoc tuning. Only if residuals persist *and* organize into a small taxonomy with strong null regimes does the framework earn continued investment.

That is the intended end state of this paper: not a new metaphysical proclamation, but a new scientific posture. Move from elegance that is enforced by hidden priors to structure that is earned by explicit gluing. Treat global models as certificates, not as foundations. Treat dark matter first as an obstruction, not as an entity. And demand, above all, functorial predictability: the ability to carry truth across contexts without losing coherence or multiplying knobs.

References

A. Category Theory, Topoi, Sheaves, Descent, and Cohomology

1. Mac Lane, S. (1971/1998). *Categories for the Working Mathematician* (2nd ed.). Springer.
2. Leinster, T. (2014). *Basic Category Theory*. Cambridge University Press.
3. Awodey, S. (2010). *Category Theory* (2nd ed.). Oxford University Press.
4. Borceux, F. (1994). *Handbook of Categorical Algebra* (Vols. 1–3). Cambridge University Press.
5. Mac Lane, S., & Moerdijk, I. (1992). *Sheaves in Geometry and Logic: A First Introduction to Topos Theory*. Springer.
6. Johnstone, P. T. (1977). *Topos Theory*. Academic Press.
7. Johnstone, P. T. (2002). *Sketches of an Elephant: A Topos Theory Compendium* (Vols. 1–2). Oxford University Press.
8. Lurie, J. (2009). *Higher Topos Theory*. Princeton University Press.
9. Bredon, G. E. (1997). *Sheaf Theory* (2nd ed.). Springer.
10. Grothendieck, A. (1971). *Revêtements étales et groupe fondamental (SGA 1)*. Springer Lecture Notes in Mathematics.
11. Vistoli, A. (2008). *Grothendieck Topologies, Fibered Categories and Descent Theory*. (Lecture notes; widely circulated).

B. Topos/Sheaf Frameworks for Quantum Theory, Contextuality, and "No Global Section" Results

12. Bell, J. S. (1964). On the Einstein Podolsky Rosen paradox. *Physics Physique Fizika*, 1, 195–200.
13. Clauser, J. F., Horne, M. A., Shimony, A., & Holt, R. A. (1969). Proposed experiment to test local hidden-variable theories. *Physical Review Letters*, 23, 880–884.
14. Kochen, S., & Specker, E. P. (1967). The problem of hidden variables in quantum mechanics. *Journal of Mathematics and Mechanics*, 17, 59–87.
15. Spekkens, R. W. (2005). Contextuality for preparations, transformations, and unsharp measurements. *Physical Review A*, 71, 052108. doi:10.1103/PhysRevA.71.052108

16. Isham, C. J., & Butterfield, J. (1998). A topos perspective on the Kochen–Specker theorem I–IV (series). *International Journal of Theoretical Physics*, 37 (and related papers).

17. Döring, A., & Isham, C. J. (2008). A topos foundation for theories of physics I: Formal languages for physics. *Journal of Mathematical Physics*, 49, 053515.

18. Döring, A., & Isham, C. J. (2008). A topos foundation for theories of physics II: Daseinisation and the spectral presheaf. *Journal of Mathematical Physics*, 49, 053516.

19. Döring, A., & Isham, C. J. (2010–2011). A topos foundation for theories of physics III–IV (series). *Journal of Mathematical Physics* (and related venues).

20. Heunen, C., Landsman, N. P., & Spitters, B. (2009). A topos for algebraic quantum theory. *Communications in Mathematical Physics*, 291, 63–110.

21. Abramsky, S., & Brandenburger, A. (2011). The sheaf-theoretic structure of non-locality and contextuality. *New Journal of Physics*, 13, 113036. doi:10.1088/1367-2630/13/11/113036

22. Abramsky, S., Barbosa, R. S., Kishida, K., Lal, R., & Mansfield, S. (2015). Contextuality, cohomology and paradox. arXiv:1502.03097.

C. Inverse Problems, Regularization, Bayesian Updating, and Model Selection

23. Hadamard, J. (1923). *Lectures on Cauchy's Problem in Linear Partial Differential Equations*. Yale University Press.

24. Tikhonov, A. N., & Arsenin, V. Y. (1977). *Solutions of Ill-Posed Problems*. Wiley.

25. Tarantola, A. (2005). *Inverse Problem Theory and Methods for Model Parameter Estimation*. SIAM.

26. Kaipio, J., & Somersalo, E. (2005). *Statistical and Computational Inverse Problems*. Springer.

27. Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90, 773–795. doi:10.1080/01621459.1995.10476572

28. Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics*, 6, 461–464.

29. Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19, 716–723.

30. MacKay, D. J. C. (2003). *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press.

D. Observational Cosmology, "Light-Cone Reconstruction," and Symmetry Priors

31. Robertson, H. P. (1935). Kinematics and world-structure. *Astrophysical Journal*, 82, 284–301.

32. Walker, A. G. (1937). On Milne's theory of world-structure. *Proceedings of the London Mathematical Society*, s2-42, 90–127.

33. Kristian, J., & Sachs, R. K. (1966). Observations in cosmology. *Astrophysical Journal*, 143, 379–399.

34. Ellis, G. F. R. (1971). Relativistic cosmology. In *Proceedings of the International School of Physics "Enrico Fermi"* (and related review literature).

35. Ellis, G. F. R., & Stoeger, W. (1987–1988). The "fitting problem" in cosmology (papers on averaging/backreaction and inference from observations).

36. Ehlers, J., Geren, P., & Sachs, R. K. (1968). Isotropic solutions of the Einstein–Liouville equations. *Journal of Mathematical Physics*, 9, 1344–1349. doi:10.1063/1.1664720

37. Noether, E. (1918). Invariante Variationsprobleme. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen*, 235–257.

E. Dark Matter Evidence and Multi-Probe Gravitational Inference

38. Zwicky, F. (1933). Die Rotverschiebung von extragalaktischen Nebeln. *Helvetica Physica Acta*, 6, 110–127.

39. Rubin, V. C., & Ford, W. K. Jr. (1970). Rotation of the Andromeda Nebula from a spectroscopic survey of emission regions. *Astrophysical Journal*, 159, 379–403.

40. Bosma, A. (1981). *21-cm Line Studies of Spiral Galaxies* (PhD thesis; foundational rotation-curve mass discrepancy work).

41. Sofue, Y., & Rubin, V. (2001). Rotation curves of spiral galaxies. *Annual Review of Astronomy and Astrophysics*, 39, 137–174.

42. Bartelmann, M., & Schneider, P. (2001). Weak gravitational lensing. *Physics Reports*, 340, 291–472. doi:10.1016/S0370-1573(00)00082-X

43. Hoekstra, H., & Jain, B. (2008). Weak gravitational lensing and its cosmological applications. *Annual Review of Nuclear and Particle Science*, 58, 99–123.

44. Markevitch, M., Gonzalez, A. H., Clowe, D., et al. (2004). Direct constraints on the dark matter self-interaction cross section from the merging cluster 1E 0657–56. *Astrophysical Journal*, 606, 819–824.

45. Clowe, D., Bradac, M., Gonzalez, A. H., et al. (2006). A direct empirical proof of the existence of dark matter. *Astrophysical Journal Letters*, 648, L109–L113.

46. Navarro, J. F., Frenk, C. S., & White, S. D. M. (1997). A universal density profile from hierarchical clustering. *Astrophysical Journal*, 490, 493–508.

47. Planck Collaboration (Aghanim, N., et al.) (2020). Planck 2018 results. VI. Cosmological parameters. *Astronomy & Astrophysics*, 641, A6.

48. Eisenstein, D. J., Zehavi, I., Hogg, D. W., et al. (2005). Detection of the baryon acoustic peak in the large-scale correlation function of SDSS luminous red galaxies. *Astrophysical Journal*, 633, 560–574.

49. Bertone, G., Hooper, D., & Silk, J. (2005). Particle dark matter: Evidence, candidates and constraints. *Physics Reports*, 405, 279–390.

50. Tulin, S., & Yu, H.-B. (2018). Dark matter self-interactions and small-scale structure. *Physics Reports*, 730, 1–57.

F. Alternatives: Modified Gravity, Varying G, and Precision Tests

51. Dirac, P. A. M. (1937). The cosmological constants. *Nature*, 139, 323.

52. Brans, C., & Dicke, R. H. (1961). Mach's principle and a relativistic theory of gravitation. *Physical Review*, 124, 925–935.

53. Milgrom, M. (1983). A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis. *Astrophysical Journal*, 270, 365–370.

54. Bekenstein, J. D. (2004). Relativistic gravitation theory for the modified Newtonian dynamics paradigm. *Physical Review D*, 70, 083509. doi:10.1103/PhysRevD.70.083509

55. Famaey, B., & McGaugh, S. S. (2012). Modified Newtonian Dynamics (MOND): Observational phenomenology and relativistic extensions. *Living Reviews in Relativity*, 15, 10. doi:10.12942/lrr-2012-10

56. Clifton, T., Ferreira, P. G., Padilla, A., & Skordis, C. (2012). Modified gravity and cosmology. *Physics Reports*, 513, 1–189. doi:10.1016/j.physrep.2012.01.001

57. Will, C. M. (2014). The confrontation between general relativity and experiment. *Living Reviews in Relativity*, 17, 4.

58. Uzan, J.-P. (2011). Varying constants, gravitation and cosmology. *Living Reviews in Relativity*, 14, 2.

G. Renormalization, Effective Field Theory, and Multi-Scale "Flow" Thinking

59. Wilson, K. G. (1971). Renormalization group and critical phenomena I: Renormalization group and the Kadanoff scaling picture. *Physical Review B*, 4, 3174–3183.

60. Wilson, K. G., & Kogut, J. (1974). The renormalization group and the epsilon expansion. *Physics Reports*, 12, 75–199.

61. Weinberg, S. (1979). Phenomenological Lagrangians. *Physica A*, 96, 327–340.

62. Polchinski, J. (1984). Renormalization and effective Lagrangians. *Nuclear Physics B*, 231, 269–295.

63. Baez, J. C., & Stay, M. (2011). Physics, topology, logic and computation: A Rosetta Stone. In *New Structures for Physics* (Springer Lecture Notes in Physics).

64. Coecke, B., & Kissinger, A. (2017). *Picturing Quantum Processes: A First Course in Quantum Theory and Diagrammatic Reasoning*. Cambridge University Press.

The author has published over 4,800 AI-assisted papers at: https://www.researchgate.net/profile/Douglas-Youvan/research . Google Scholar has indexed these preprints, and if you want a particular reference, the AI mode of a Google search (Gemini) has efficiently catalogued these preprints.

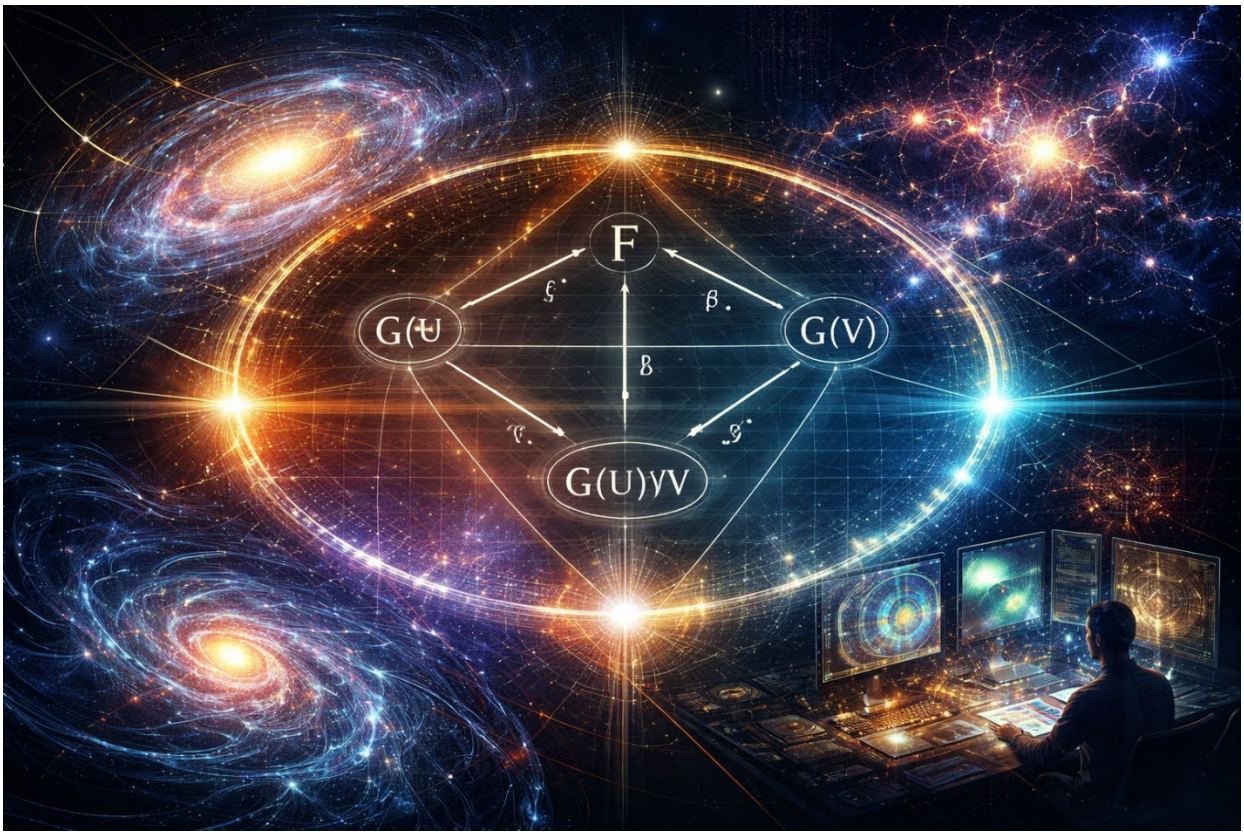
F
G(U)
G(V)
G(U)∨V